

Aggregate Consequences of Relationship Dynamics in Trade*

Pulak Ghosh

IIM, Bangalore

Edward Wiles

Harvard University

Please **click here** for the latest version.

27 October 2026

Abstract

Most trade takes place within long-term buyer-seller relationships whose value reflects past experience and, when contract enforcement is limited, the expectation of future business. We study how these relationship dynamics shape aggregate adjustment to changes in trade costs. We develop a dynamic multi-location, multi-sector Ricardian model with relationship-specific learning and relational contracts. Accumulated relationship value slows adjustment as some firms retain incumbent suppliers despite cheaper alternatives, while improved outside options can offset this by undermining the incentives needed to sustain those relationships. A property of the optimal contract that we call scale equivariance permits tractable aggregation of these micro-level dynamics. Using rich firm-to-firm transaction data from India, we establish that proportional growth within relationships is remarkably similar across firm characteristics and experience—supportive of the scale equivariance property—and varies with contract enforcement. We use these empirical results as moments to estimate the model’s learning and contracting parameters, and study a reduction in interstate trade costs calibrated to a major Indian domestic integration reform. Both dynamic mechanisms are quantitatively important: it takes 15 quarters for the economy to reach 90% of its long-run trade elasticity, compared with 30 quarters when relationships can continue even if their contracts violate the post-shock incentive constraint.

*We are grateful to Pol Antràs, David Atkin, Abhijit Banerjee, Arnaud Costinot, Dave Donaldson, Esther Dufo, Elhanan Helpman, Marc Melitz, Danny O’Connor and attendees of various MIT and Harvard internal seminars for helpful comments. Corresponding author: edward.m.wiles@gmail.com.

1 Introduction

Most trade takes place within long-term buyer-seller relationships that are valuable because of both past experience and the expectation of future business. These relationship dynamics shape how firms respond to changes in trade costs. First, a buyer that has accumulated years of experience with its supplier may prefer to retain this supplier despite cheaper alternatives becoming available, and those that do switch must build new relationships from scratch. Second, even among firms that prefer to retain their incumbent supplier, the presence of cheaper alternatives may itself weaken the implicit promise of future business needed to sustain trade when contracts are hard to enforce. Yet, despite growing evidence on the importance of these micro-level relationship dynamics, we know comparatively little about how they shape the aggregate response to changes in trade costs, in part because it is challenging to incorporate them into workhorse quantitative trade models.

In this paper, we develop a quantitative trade model that incorporates relationship dynamics, establish new empirical facts using rich microdata from India about the scope of relationship dynamics and how they relate to contract enforcement, and quantify how these relationships shape the aggregate response to changes in trade costs. The model features relationship-specific learning that generates gradual buildup over the course of a relationship, and relational contracts in which the prospect of continued business provides incentives to honor current obligations. We show that these features can be incorporated tractably into a multi-location, multi-sector general equilibrium Ricardian model due to a property that we call scale equivariance. We use data covering the universe of domestic firm-to-firm sales over nine years for all sellers registered in a large Indian state to show that the proportional buildup of trade is remarkably similar across firm characteristics and experience, but varies with the quality of the destination state’s contract enforcement. Finally, we estimate the model using the microdata and study the transition path following a reduction in interstate trade costs calibrated to a major Indian domestic integration reform.

In the first part of the paper, we develop a dynamic Ricardian trade model that incorporates relationship-specific learning and optimal relational contracts between buyers and sellers. The aggregate environment follows the standard multi-location, multi-sector framework of [Eaton and Kortum \(2002\)](#) and [Caliendo and Parro \(2015\)](#), in which each location has a Frechet-distributed productivity draw for every good and production combines labor and intermediate inputs. We depart from this framework by introducing specialized importers that, for each good, choose which origin to source from and then match with an individual supplier in the chosen origin. Trade therefore takes place within buyer-seller relationships rather than spot transactions.

Within each relationship, the buyer and seller face two frictions. First, the match has some probability of being commercially unsuitable, in which case any given shipment may fail to meet the buyer’s requirements with an exogenous probability. Both parties learn about this pair-specific suitability through repeat transactions and update their beliefs using Bayes’ Rule. As successful

transactions accumulate, beliefs of commercial suitability improve, and the buyer wants to increase their order size. Relationships therefore create *history-dependence* that will affect how the buyer responds when trade costs change. Second, there is a hold-up problem: after the seller has incurred production and shipment costs, the buyer can attempt to avoid payment, in which case the seller can only recover a fraction through courts. To sustain payment in equilibrium, the contract must satisfy a Dynamic Incentive Compatibility Constraint (DICC): the future value of the relationship must exceed the one-off gain from hold-up. The optimal order size is therefore the minimum of the quantity implied by current beliefs about suitability and the quantity that the DICC can sustain. Relationships therefore also create *future-dependence*: if changes in trade costs improve the outside option, they tighten the DICC today.

These frictions jointly generate relationship dynamics. Learning causes quantity to grow as a relationship ages and makes older relationships more valuable. Imperfect contract enforcement can amplify this buildup—when the quantity that is feasible under the DICC is less than the quantity implied by current beliefs—particularly in early stages when the relationship is less valuable. The extent to which the DICC binds depends on the strength of contract enforcement. Importers therefore choose the origin to source from based on the expected discounted profits from these relationships, rather than simply choosing the origin with the lowest marginal cost. Relationships break up from time to time, when realizations reveal commercial unsuitability or when the good exogenously becomes obsolete, which creates a cross-sectional distribution of relationship ages.

A challenge with embedding these micro-level features into a quantitative trade model is that heterogeneous productivity would appear to imply that we need to solve a separate contracting problem for every variety. We show that this can be avoided through a property of the optimal contract that we call *scale equivariance*. While productivity affects the absolute value traded, the ratio of trade early in a relationship to trade late in a relationship is the same for all varieties. This property follows from CES demand and both production costs and the DICC being homogeneous in productivity. Thus, it is sufficient to solve a normalized contracting problem. Scale equivariance continues to hold when aggregate conditions change: normalized contracts depend on future changes in wages, input prices, and demand, but not on their absolute levels. We also show that the same scaling logic applies to other dynamic microeconomic mechanisms, including models of learning about demand, customer capital accumulation, and investment in physical capital.

Combining scale equivariance with Frechet-distributed productivity allows us to aggregate the micro-level relationship dynamics tractably. In steady state, bilateral trade flows follow a standard gravity equation with two aggregate wedges: an extensive margin wedge that summarizes how differences across locations in learning and contract enforcement distort the choice of origin, and an intensive margin wedge capturing the average quantity distortion across relationships of different ages. Despite the rich microeconomic dynamics, the long-run trade elasticity remains constant and equal to the Frechet shape parameter, and comparisons between steady states can be calculated using the standard exact hat algebra method of [Dekle et al. \(2008\)](#).

The transition path to the new steady state following a change in trade costs highlights the two dynamic implications of relationships. First, importers may retain incumbent suppliers despite now-cheaper alternatives because of learning-induced relationship-specific value, while those that switch must build new relationships from scratch. This slows the reallocation of trade. Second, when an alternative origin becomes cheaper, the value of an incumbent relationship falls relative to the outside option. This tightens the DICC, which reduces contractible quantity and can cause the relationship to break down altogether. This can partly offset the learning-induced stickiness, as well as generate anticipatory responses to future changes in economic conditions.

We show that the transition path can be computed in relative changes. The additional complication is that the decision to retain an incumbent supplier, and the contract governing that retention, depend on the ratio of the value of the incumbent to the value of starting fresh elsewhere. An additional property of the Frechet distribution makes this tractable: this ratio follows a Lomax distribution that depends only on the pre-shock trade shares and relative changes in aggregate variables. Together with scale equivariance, this permits tractable aggregation and computation.

In the second part of the paper, we use rich microdata from India to discipline the model and to establish two new empirical facts about relationship dynamics. Our data covers the near-universe of domestic firm-to-firm goods shipments involving firms registered for value-added taxes in a large Indian state, including sales and purchases from other states, from 2018 through 2026. The data includes quantity and product codes for all items in each shipment. Consistent tax identifiers allow us to track buyer-seller relationships over an eight year period and to examine how their development varies with the characteristics of both parties.

We first measure how value traded grows with relationship age. While the importance of within-relationship growth has been established elsewhere, we build on the literature by deriving our estimating equation directly from the model, which makes clear the necessary set of fixed effects and shows that the coefficients can be interpreted as structural wedges. Value traded in a relationship starts at around half of its long-run level, expands gradually over around eight quarters, and becomes stationary thereafter. This growth is driven almost entirely by quantity, with little change in unit prices. We also show that the probability that a given relationship trades again is strongly increasing in relationship age, eventually reaching a mature survival rate of around 95%.

We then present two new empirical facts about relationship dynamics. First, to assess the empirical plausibility of scale equivariance and to distinguish between alternative explanations for relationship buildup, we test whether the value growth curve varies systematically with firm characteristics. We find that the (proportional) buildup of trade within a relationship is remarkably similar across buyer and seller size, age, and location- and product-experience, despite large differences in absolute trade values. This provides empirical support for the separability between the scale of trade and relationship dynamics that plays an important role in the model. It also weighs against explanations based primarily on transferable learning about locations or products, which would predict less growth among experienced firms.

Second, we study how this within-relationship value growth curve varies with contract enforcement. We exploit differences in court congestion across destination states, following [Boehm and Oberfield \(2020\)](#) in using the age of the relevant high court as an instrument. We find that early transactions are closer to their mature level when enforcement is stronger, with the predicted curve at 2 standard deviations of enforcement beyond the mean reaching maturity after around one year (as compared to around two years at the mean enforcement level). This effect is stronger in more relationship-specific sectors. This is consistent with the model’s prediction that stronger enforcement relaxes the DICC and permits greater trade earlier in a relationship. Moreover, it provides direct reduced form evidence as to one mechanism that causes relationships to start small.

In the third part of the paper, we use the value growth curve and its variation across enforcement environments to jointly estimate the learning and contracting parameters. We assume common learning parameters across origins, destinations, and sectors, and allow enforcement to vary by destination through a two-parameter logistic relationship with court congestion. This uses the common pattern of relationship growth to discipline learning, while its variation with enforcement identifies the contracting friction.

We combine our estimates with a new India-wide matrix of aggregate bilateral trade flows, an input-output matrix, and household expenditure data to calibrate an economy with 20 locations and 15 sectors. We then study a permanent reduction in interstate trade costs calibrated to estimates from ([Barnwal et al., 2026](#)) as to the impact of a major 2017 domestic integration reform on trucking delays at state borders. To isolate the role of changes in future surplus, we compare the transition path under the full model with a version in which we hold the normalized relationship growth wedges fixed at their pre-shock levels and allow for violations of the post-shock DICC.

We find that relationship dynamics have important effects on the speed of trade adjustment. The trade elasticity on impact is 2.85, between the intensive margin elasticity of 2 and the long-run elasticity of 4. Among firms that would switch where they source from in a static model, 67% instead retain their incumbent supplier. The model therefore endogenously generates a meaningfully slower short-run response, consistent with recent evidence that short-run elasticities are lower than long-run elasticities ([Boehm et al., 2023](#)).

At the same time, the loss of surplus in incumbent relationships provides an offsetting force. When we hold the relationship growth wedges fixed at their pre-shock values, the impact elasticity is instead 2.52, and 88% of firms retain their incumbent suppliers: relational contracting is no longer sustainable for many otherwise profitable relationships. The fixed-wedge version takes 30 quarters to reach 90% of the long-run trade elasticity, compared to 15 quarters under the full model. Thus, different microeconomic mechanisms generate meaningfully different transition paths.

This paper contributes to three literatures. First, a large literature studies the gradual expansion of exporters into new markets, emphasizing sunk costs, learning, and customer accumulation ([Das et al., 2007](#); [Berman et al., 2019](#); [Fitzgerald et al., 2024](#)), while a related empirical literature docu-

ments that trade also builds gradually within individual buyer-seller relationships (Monarch and Schmidt-Eisenlohr, 2023; Heise, 2024; Brugués, 2026). We contribute to this evidence by showing that relationship growth is remarkably similar across firm characteristics but varies with contract enforcement, and by deriving our regression specification directly from a model, which informs the choice of fixed effects and shows that the estimated coefficients can be interpreted as structural wedges.¹ On the theoretical side, recent work has developed rich models of buyer-seller relationships that feature learning (Eaton et al., 2025; Kohn et al., 2026), matching (Fontaine et al., 2023; Grossman et al., 2024), and relational contracting (Boehm et al., 2024). We extend this literature by embedding learning and relational contracts into a multi-location, multi-sector quantitative trade model, allowing us to study how relationship dynamics interact with general equilibrium changes in wages, input prices, and trade flows along the transition path following a trade shock.

Second, we connect the literature on contracting frictions in trade (Antràs, 2003; Nunn, 2007; Antras and Foley, 2015; Boehm and Oberfield, 2020; Chor and Ma, 2025) with the empirical literature on relational contracting (Macchiavello and Morjaria, 2015; Blouin and Macchiavello, 2019; Macchiavello and Morjaria, 2021; Harris and Nguyen, 2025; Wiles and Houeix, 2026). The former emphasizes how imperfect enforcement shapes trade patterns and production decisions, typically in static settings; the latter shows the importance of future rents in sustaining trade when enforcement is limited. We incorporate this well-established dynamic mechanism into a workhorse quantitative trade model, allowing us to study how changes in future relationship surplus affect aggregate trade adjustment.

Third, we contribute to the growing literature on dynamic quantitative trade models (Caliendo et al., 2019; Kleinman et al., 2023; Chen et al., 2026; Kleinman et al., 2026), which has studied dynamic adjustment through labor mobility, capital accumulation, and supplier choice. We show that the gradual buildup and future dependence of buyer-seller relationships, which have been documented extensively in microeconomic settings, has important implications for the speed of aggregate trade adjustment.² Methodologically, we identify a property that we call scale equivariance that allows a broad class of microeconomic mechanisms to be incorporated into quantitative trade models with transition paths computed in relative changes.

Section 2 describes the model. Section 3 presents the empirical analysis. Section 4 uses the empirical evidence to estimate the model. Section 5 reports the quantitative analysis. Section 6 concludes.

¹Araujo et al. (2016) find that the rate of growth of firms’ exports to a new country are negatively correlated with country-level measures of contract enforcement. In contrast, we study growth within individual buyer-seller relationships and use an instrument for contract enforcement.

²Perhaps the closest paper to ours is Chen et al. (2026), who develop a Ricardian model where importers can only switch to a new supplier when a Calvo-style shock occurs, thereby generating slow adjustment to changes in trade costs. Our model features (endogenous) within-relationship dynamics, which generates endogenous switching as a function of accumulated within-relationship value. We also include forward-looking relational contracting.

2 Model

In this section, we develop a quantitative general equilibrium Ricardian model featuring relationship dynamics. The aggregate environment follows the standard multi-location, multi-sector framework of [Eaton and Kortum \(2002\)](#) and [Caliendo and Parro \(2015\)](#). We introduce a micro-level environment that features relationship-specific learning and optimal relational contracts, and show that aggregation remains tractable due to a property that we call scale equivariance.

2.1 Environment

There are D locations, indexed o for origin and d for destination, and S sectors, indexed j or k . Time t is discrete with common discount factor β .

2.1.1 Sectoral Composites

In each destination d , there is a competitive CES aggregator associated with each sector j that combines a unit mass of sector j varieties with elasticity $\sigma_j > 1$ to produce a sectoral composite good $Q_{d,j,t}$ with CES price index $P_{d,j,t}$:

$$Q_{d,j,t} = \left(\int_0^1 q_{d,j,t}(\omega)^{\frac{\sigma_j-1}{\sigma_j}} d\omega \right)^{\frac{\sigma_j}{\sigma_j-1}}, \quad P_{d,j,t} = \left(\int_0^1 p_{d,j,t}(\omega)^{1-\sigma_j} d\omega \right)^{\frac{1}{1-\sigma_j}}, \quad (1)$$

where $q_{d,j,t}(\omega)$ and $p_{d,j,t}(\omega)$ represent the quantity of variety ω that the aggregator purchases and the price that it pays. The aggregator purchases varieties from specialized importers, discussed in Section 2.1.4. These composite goods are used both in final consumption and as intermediates in downstream production of individual varieties. Define $E_{d,j,t} \equiv P_{d,j,t} Q_{d,j,t}$ as total expenditure on sector j composites in destination d , and $\nu_{d,j,t} \equiv E_{d,j,t}^{1/\sigma_j} P_{d,j,t}^{(\sigma_j-1)/\sigma_j}$ as a destination-sector demand shifter. Then, cost minimization implies that the aggregator's inverse demand curve for a particular variety ω is

$$p_{d,j,t}(\omega) = \nu_{d,j,t} q_{d,j,t}(\omega)^{-1/\sigma_j}. \quad (2)$$

2.1.2 Production

Consider a given variety ω in sector j . Every origin o is capable of producing this variety, with idiosyncratic origin-sector productivity $z_{o,j}(\omega)$ drawn from $\text{Frechet}(1, \theta_j)$. Production combines labor and the composite goods in (1) from each sector used as intermediates with a Cobb-Douglas production function. The marginal cost for origin o to produce and deliver a sector- j good ω to destination d , denoted $c_{od,j,t}(\omega)$, is given by

$$c_{od,j,t}(\omega) = \frac{\tau_{od,j,t}}{z_{o,j}(\omega)} \left(w_{o,t}^{\gamma_j} \prod_{k=1}^S (P_{o,k,t})^{\gamma_{k,j}} \right), \quad (3)$$

where τ_{od} are bilateral cost shifters, w_o is the wage in o , $\gamma_{k,j}$ is the cost share of the sector- k composite good used as an intermediate in the production of sector j , and $\gamma_j = 1 - \sum_k \gamma_{k,j}$ is the labor cost share. We will often refer to $c_{o,j,t} \equiv \left(w_{o,t}^{\gamma_j} \prod_{k=1}^S (P_{o,k,t})^{\gamma_{k,j}} \right)$ as the input bundle unit cost and $\kappa_{od,j,t} \equiv \tau_{od,j,t} c_{o,j,t}$ as the component of marginal cost common across goods. Marginal cost can therefore be written as $c_{od,j,t}(\omega) = \kappa_{od,j,t} / z_{o,j}(\omega)$.

2.1.3 Households

There is a mass L_d of households in each destination d who supply labor and constitute final demand. Households have Cobb-Douglas utility over the sectoral composite goods with expenditure share on sector j denoted $\alpha_{d,j}$. Formally, let $C_{d,j,t}$ denote consumption of the sector- j composite in destination d . Total consumption and the consumer price index are therefore given by

$$C_{d,t} = \prod_{j=1}^S \left(\frac{C_{d,j,t}}{\alpha_{d,j}} \right)^{\alpha_{d,j}}, \quad P_{d,t}^C = \prod_j P_{d,j,t}^{\alpha_{d,j}}. \quad (4)$$

Final expenditure satisfies $P_{d,j,t} C_{d,j,t} = \alpha_{d,j} I_{d,t}$, where $I_{d,t}$ is the consumer's income. Any profits are rebated lump-sum to consumers in the location of the firm that earns them.

2.1.4 Importers

For every destination-sector-variety (d, j, ω) , there is a unique importer. The importer selects the origin o to source the good from that maximizes expected discounted profits—described in more detail shortly—and sells it to the relevant sectoral aggregator at price $p_{d,j,t}(\omega)$. Importers compete monopolistically across ω , taking the common demand shifter $\nu_{d,j,t}$ as given, and earn revenue $\nu_{d,j,t} q_{d,j,t}(\omega)^{(\sigma_j - 1)/\sigma_j}$.

2.1.5 Supplier Relationships

An importer that has selected an origin o is matched at random with an individual supplier in o . Within each origin-sector-variety (o, j, ω) , there is an infinite pool of *ex ante* identical suppliers, all of whom have access to the production technology in (3). We refer to the number of periods that a match has transacted as the age of the relationship, $a \in \{0, 1, 2, \dots\}$. Relationships can end in two ways. First, at the end of any period, either party can terminate the relationship. If this happens, the importer matches with a new supplier in o . Second, at the end of every period, a share $1 - \delta$ of goods become obsolete and are replaced by new goods. If this happens, a new vector of productivities is drawn at the start of the next period, $\{z_{o,j}(\omega)\}_o$, and the importer selects an origin o to source from at the start of the next period.

2.2 Contracting

After matching, a given importer-supplier (“buyer-seller”) pair play a dynamic contracting game. Throughout this subsection, we suppress indices for origin, destination, sector, and variety (these enter through the marginal cost $c_{odj,t}(\omega)$ and the demand shifter ν_{djt}).

A given match is commercially unsuitable with probability μ , independent of all other characteristics. If a match is unsuitable, after delivery in each period the good becomes worthless with probability $1 - \lambda$ (it is incompatible with the CES aggregator’s technology); if the match is suitable, the corresponding probability is 0. Neither party can observe directly whether a match is suitable, but both update their beliefs using Bayes’ Rule based on prior realizations within the relationship. After a successful realizations, the posterior probability that the match is unsuitable is

$$\mu_a = \frac{\mu\lambda^a}{1 - \mu + \mu\lambda^a}, \quad (5)$$

and the probability of a successful transaction is $p_a \equiv 1 - \mu_a + \mu_a\lambda$. All else equal, relationships with more successful prior realizations are more valuable, with $\mu_a \rightarrow 0$ and $p_a \rightarrow 1$ as $a \rightarrow \infty$. As it will be optimal for the buyer to exit the relationship if a bad realization occurs, age is a sufficient state for suitability.

Within each period, the parties face a standard hold-up problem. First, the buyer orders quantity $q_{a,t}$ and promises payment $x_{a,t}$. Second, the seller produces and ships the good at cost $c_t q_{a,t}$. We assume that production and shipment are verifiable before the goods have arrived (e.g., through a third-party logistics provider), but that commercial suitability is not. Third, after production and shipment are sunk, the buyer can attempt to renegotiate the promised payment downwards, which for simplicity we model as attempted nonpayment. If the buyer does so, the seller can recover a fraction $1 - \xi$ through the legal process, where $\xi \in [0, 1]$ captures the weakness of formal contract enforcement, which may vary by origin, destination, or sector. Finally, the realization occurs.

We consider Perfect Public Equilibria supported by grim-trigger strategies: the relationship terminates permanently if the buyer attempts nonpayment. Let $V_{a,t}^b$ denote the buyer’s continuation value in a relationship of age a in period t , given by

$$V_{a,t}^b = p_a \nu_t q_{a,t}^{(\sigma-1)/\sigma} - x_{a,t} + \beta \delta \left(p_a V_{a+1,t+1}^b + (1 - p_a) V_{t+1}^{b,\text{out}} \right), \quad (6)$$

where $V_{t+1}^{b,\text{out}}$ is the buyer’s outside option. The outside option is given by the value of starting a new relationship, which may be in the same origin or a different origin. In any equilibrium in which the buyer repays, a Dynamic Incentive Compatibility Condition (DICC) must hold:

$$\xi x_{a,t} \leq \beta \delta p_a (V_{a+1,t+1}^b - V_{t+1}^{b,\text{out}}). \quad (7)$$

Intuitively, the left hand side is the one-time gain the buyer can earn from nonpayment, and the right hand side is the expected future value gained from preserving this relationship into age $a + 1$ relative to starting fresh with a new supplier. If the outside option is to remain in the same origin, $V_{t+1}^{b,\text{out}} = V_{0,t+1}^b$. The DICC shows that strong formal enforcement (low ξ) and high future rents (high $V_{a+1,t+1}^b$) are substitutes in sustaining trade.

Optimal Contract The buyer makes a take-it-or-leave-it offer at the start of the match, which is a sequence $(q_a, x_a)_{a=0}^\infty$, that maximizes their expected discounted profits subject to satisfying (7) and the seller at least breaking even ($x_{a,t} \geq c_t q_{a,t}$) every period.³ It is straightforward to observe that the optimal payment satisfies $x_{a,t} = c_t q_{a,t}$: paying the supplier more would both lower the buyer's discounted profits and tighten the DICC. The contracting problem therefore reduces to choosing quantities. Absent the DICC, the buyer would choose $q_a^{\text{learn}} = (((\sigma - 1)/\sigma) \cdot p_a \nu / c)^\sigma$. The DICC, however, imposes an upper bound on contractible quantity. Combining the DICC with $x_{a,t} = c_t q_{a,t}$ therefore yields

$$q_{a,t} = \min \left\{ \underbrace{\left(\frac{\sigma - 1}{\sigma} \frac{p_a \nu_t}{c_t} \right)^\sigma}_{q_{a,t}^{\text{learn}}}, \underbrace{\frac{\beta \delta p_a (V_{a+1,t+1}^b - V_{t+1}^{b,\text{out}})}{\xi c_t}}_{q_{a,t}^{\text{DICC}}} \right\}. \quad (8)$$

Equation (8) shows that trade is distorted downwards by the two frictions. First, uncertainty about commercial suitability ($p_a < 1$) lowers desired quantity. Second, weak contract enforcement ($\xi > 0$) means that quantity may be constrained by the surplus available in the relationship.

For a given value of the outside option, $\tilde{V}_{t+1}^{b,\text{out}}$, the solution can be obtained recursively using (6) and (8). The terminal condition is the limiting case where $p_a = 1$. In this limiting case, if the continuation value from playing q_∞^{learn} every period satisfies the DICC, this is the solution; otherwise, it is the largest value of q such that playing q every period satisfies the DICC. If the contracting friction is sufficiently severe, or the outside option is sufficiently high, there may be no such q , in which case trade is not feasible. Indeed, long-run contractible quantity does not smoothly decline all the way from first-best to zero as one increases the outside option: it declines initially, but at some point drops discontinuously to zero. Intuitively, reducing (stationary) quantity lowers the temptation but also lowers the discounted sum of future profits, while the outside option remains fixed. Importantly, this can occur even though the buyer prefers this relationship to their outside option: it may be that $\xi x_{\infty,t+1}/(\beta \delta) > V_{\infty,t+1}^b - V_{t+1}^{b,\text{out}} > 0$. A strong enough outside option can therefore cause a relationship to break down even though both parties would prefer to keep trading. In the case where the environment is stationary and the outside option is to remain in the same origin, equilibrium is the fixed point at which the assumed value of the outside option,

³We let the buyer choose the contract as allowing the seller to do so would cause the price to decline substantially with relationship age (as it is optimal for incentives to backload the buyer's rents), which is contrary to our finding in the data of a near flat price profile.

$\tilde{V}^{b,\text{out}}$, equals the value of a new relationship implied by backwards induction, V_0^b .

2.3 Scale Equivariance

While the solution to the contracting model described in Section 2.2 must be computed numerically, it has an important property that substantially simplifies incorporating it into a quantitative trade model: optimal contracts are *scale equivariant*.⁴ In particular, while varieties differ by productivity $z_{o,j}(\omega)$ and fundamentals $\nu_{d,j,t}$ and $\kappa_{od,j,t}$, these differences affect the absolute scale of quantities and value traded but not the proportional dynamics of a relationship. This means that we need only solve a suitably normalized contracting problem and then scale the normalized solution, rather solve separate problems for all possible values of productivity and fundamentals. This allows us to separate the microeconomic contracting problem from the underlying productivity heterogeneity and general equilibrium variables that characterize quantitative trade models.

Recall that the marginal cost of supplying variety ω from origin o to destination d in sector j is $c_{od,j,t}(\omega) = \kappa_{od,j,t}/z_{o,j}(\omega)$. The following proposition establishes that the solutions to the contracting problem can be expressed entirely as a frictionless solution multiplied by age-specific shifters.

Proposition 1 (Scale Equivariance of Relational Contracts). *Consider the contracting problem described in Section 2.2. Assume the buyer's outside option is always to re-match with a new supplier in origin o . For every calendar date t and every relationship $a \in \{0, 1, 2, \dots\}$, consider the optimal continuation contract for a relationship entering period t with beliefs μ_a . There exist scalars $\phi_{a,odj,t}^q, \phi_{a,odj,t}^x, \phi_{a,odj,t}^V$ such that*

$$\begin{aligned} q_{a,odj,t}(\omega) &= z_{o,j}(\omega)^\sigma \nu_{dj,t}^\sigma \kappa_{odj,t}^{-\sigma} \phi_{a,odj,t}^q \\ x_{a,odj,t}(\omega) &= z_{o,j}(\omega)^{\sigma-1} \nu_{dj,t}^\sigma \kappa_{odj,t}^{1-\sigma} \phi_{a,odj,t}^x \\ V_{a,odj,t}^b(\omega) &= z_{o,j}(\omega)^{\sigma-1} \nu_{dj,t}^\sigma \kappa_{odj,t}^{1-\sigma} \phi_{a,odj,t}^V. \end{aligned}$$

The scalars $\phi_{a,odj,t}^q, \phi_{a,odj,t}^x, \phi_{a,odj,t}^V$ depend on odj only to the extent that (μ, λ, ξ) vary with odj , and depend on t only through the anticipated path of aggregate conditions $\nu_{d,j,t}$ and $\kappa_{od,j,t}$ expressed in relative changes. We write $\phi_{\infty,odj,t}^q, \phi_{\infty,odj,t}^x, \phi_{\infty,odj,t}^V$ as the limiting cases if posteriors beliefs assign probability 1 to commercial suitability.

All proofs are available in Appendix A. This proposition states it is sufficient to solve a suitably normalized version of the contracting problem that depends only on the contracting and learning parameters, relationship age, and the anticipated future evolution of aggregate economic conditions. Crucially, the solution does not depend on the absolute levels of productivity, demand, or costs.

⁴A mapping is said to be equivariant if transforming the inputs and applying the mapping produces the same result as applying the mapping and transforming the outputs. It is a generalization of homogeneity that allows transformations to act differently on different inputs and outputs, including scaling different variables (quantities and payments in our case) by different powers of a common factor.

The intuition is straightforward. CES revenue is homothetic, production costs are linear, and the constraints do not depend on scale. Increasing productivity by a factor $k > 0$ therefore increases $x_{a,t}$ and V_a^b by a factor $k^{\sigma-1}$, which leaves (7) and the supplier breakeven constraint unchanged and scales the objective function by $k^{\sigma-1}$, which implies that the transformed contract remains feasible and optimal.⁵ More productive varieties trade larger quantities and values, but, after normalizing by their productivity, the contracting policies are identical. In a stationary environment, these normalized contracting policies depend only on age (the sufficient state) and any *odj*-specific learning and contracting parameters. We explore this prediction empirically in Section 3, where the richness of the microdata allows us to examine whether the proportional build-up varies systematically with several important aspects of firm heterogeneity.

Scale equivariance continues to hold even when aggregate conditions evolve over time. To see this, suppose $\nu_{odj,t}$ and $\kappa_{odj,t}$ vary over time, and consider the DICC between calendar periods t and $t + 1$. Substituting the solutions from Proposition 1 into (7) gives

$$\xi \phi_{a,odj,t}^x \leq \underbrace{\left(\frac{\nu_{dj,t+1}}{\nu_{dj,t}} \right)^\sigma}_{\hat{\nu}_{d,j,t}^\sigma} \underbrace{\left(\frac{\kappa_{odj,t+1}}{\kappa_{odj,t}} \right)^{1-\sigma}}_{\hat{\kappa}_{odj,t}^{1-\sigma}} \beta \delta p_a (\phi_{a+1,odj,t+1}^V - \phi_{0,odj,t+1}^V). \quad (9)$$

The same substitutions can be done in the objective function and the seller's breakeven constraint. The normalized contracting policies, $\{(\phi_{a,odj,t}^q, \phi_{a,odj,t}^x, \phi_{a,odj,t}^V)\}$ therefore depend on aggregate economic conditions only through their relative changes over time. This property is important for the analysis of trade shocks, as general equilibrium forces will cause future wages, prices, and demand shifters to change. Indeed, because continuation values are forward-looking, the contracting solution will generally depend on the entire future path of these aggregate economic conditions. Nevertheless, the normalized policies can be computed without needing to know the absolute levels of these variables, which means they are amenable to the standard “exact hat algebra” approach of Dekle et al. (2008).

Beyond relational contracts. The tractability provided by scale equivariance is not specific to the contracting model studied here. The key property is that differences in productivity or economic conditions change the scale of a microeconomic problem but not the trade-offs governing optimal behavior once the problem is suitably normalized. This property arises in a variety of dynamic models with state variables used in trade, including models of learning about demand (Berman et al., 2019), acquiring customer capital (Fitzgerald et al., 2024), dynamic screening (Brugués, 2026), learning about counterparty types (Monarch and Schmidt-Eisenlohr, 2023), inventory manage-

⁵Perhaps the most intuitive way to see this is to inspect Equation (8). Consider doubling both c_t and ν_t for all t , and conjecture that the continuation value, $V_{a,t}^b$, also doubles at every age. It can then immediately be seen that the optimal schedule of $q_{a,t}$ is unchanged. As $x_{a,t} = c q_{a,t}$, and $V_{a,t}^b$ is simply the expected infinite sum of the stage payoff $p_a \nu_t q_{a,t}^{(\sigma-1)/\sigma} - c_t q_{a,t}$, this then implies that the stage payoff doubles for all ages. As the state transition probabilities are unchanged, it then follows that the conjecture that $V_{a,t}^b$ doubles for all ages is verified.

ment (Chen, 2025), and physical capital accumulation. Although the economic mechanisms differ substantially, the microeconomic problems can be expressed in terms of normalized policies that depend on the relevant state—such as beliefs, capital, or inventories—rather than separately on productivity or aggregate economic conditions.

The broader implication is that a number of microeconomic mechanisms can be incorporated into quantitative trade models without the need to track the joint distribution of productivity and state variables of the microeconomic problem. When trade flows at the micro-level are proportional to a common productivity term, aggregate trade depends only on the (productivity-weighted) distribution of the state variables. If the evolution of the state variables does not depend on productivity, and the extensive margin origin-choice decision admits tractable aggregation, then the resulting model can be solved in relative changes.⁶ We discuss this more general implication in Appendix B.

2.3.1 Aggregation and Market Clearing

Let $a_{d,j,t}(\omega)$ denote the age of the relationship for which d sources the sector- j good ω at time t , and let $o_{d,j,t}(\omega)$ denote the origin that it sources the good from. For ease of exposition, we suppress subscripts and simply write these as $a(\omega)$ and $o(\omega)$. The contracting problem determines the quantity produced by the supplier and ordered by the importer, $q_{a(\omega),o(\omega)d,j,t}(\omega)$, and the payment made from the importer to the supplier, $x_{a(\omega),o(\omega)d,j,t}(\omega)$. As the importer internalizes the demand curve of the CES aggregator, the price at which the importer sells to the CES aggregator is given by substituting $q_{a(\omega),o(\omega)d,j,t}(\omega)$ into the inverse demand curve in (2); if the transaction is unsuccessful (which occurs with probability $1 - p_a$ in a relationship of age a), we adopt the convention that the price is $+\infty$ and the corresponding quantity that enters (1) is zero.

Aggregate bilateral trade flows are the sum of payments over all goods, and are thus given by

$$X_{odj,t} \equiv \int_0^1 x_{a(\omega),odj,t}(\omega) 1\{o_{dj,t}(\omega) = o\} d\omega. \quad (10)$$

Similarly, let $c_{a,odj,t}(\omega)$ denote the total cost incurred by a supplier in o of good ω selling to destination d in a relationship of age a . Then, total expenditure by producers is given by

$$Y_{oj,t} \equiv \sum_d \int_0^1 c_{a_{dj,t}(\omega),odj,t}(\omega) 1\{o_{odj,t}(\omega) = o\} d\omega. \quad (11)$$

⁶One might suppose that scale equivariance generically fails in the presence of fixed costs. This is not the case. The result in Proposition 1 that the continuation value scales is stronger than needed in general because the extensive margin structure that we use in this paper is based on the Frechet distribution, which requires multiplicative scaling of value functions for aggregation. However, with the Melitz (2003)-Pareto extensive margin structure typically associated with fixed costs, it is sufficient that the post-entry operating problem (that determines the intensive margin, $x_{s,odj,t}(\omega)$) is scale equivariant. Various forms of fixed costs, such as market-entry costs, search costs, or customer setup costs that determine cutoffs can all be incorporated and aggregated analytically as in the static model.

We assume that trade is balanced, so labor market clearing in each origin implies

$$w_{o,t}L_o = \sum_j \gamma_j Y_{oj,t}. \quad (12)$$

2.4 Steady State

We first study how the dynamic contracting friction affects trade in steady state. We begin by formally defining a stationary equilibrium. The exogenous fundamentals of the economy are bilateral trade costs $\{\tau_{odj}\}$ and labor endowments $\{L_o\}$. The aggregate parameters are the labor shares $\{\gamma_j\}$, input-output cost shares $\{\gamma_{k,j}\}$, household expenditure shares $\{\alpha_{d,j}\}$, discount factor β , exogenous non-obsolescence probability δ , elasticity of substitution in the sectoral CES aggregators $\{\sigma_j\}$, and Frechet shape parameters $\{\theta_j\}$. The contracting parameters are the probability of commercial unsuitability $\{\mu_{odj}\}$, the probability of a successful transaction if a match is commercially unsuitable $\{\lambda_{odj}\}$, and the fraction of an invoice that cannot be recovered through courts $\{\xi_{odj}\}$. We collect these parameters in Θ . The endogenous aggregate variables are wages $\{w_o\}$, sectoral prices $\{P_{d,j}\}$, and sectoral expenditures $\{E_{d,j}\}$.

Definition 1 (Stationary Equilibrium). *Given exogenous fundamentals and parameters, $(\{\tau_{odj}\}, \{L_o\}, \Theta)$, a stationary equilibrium consists of aggregate endogenous variables $(\{w_o\}, \{P_{d,j}\}, \{E_{d,j}\})$, optimal contracts, sourcing decisions, and a stationary distribution of relationship ages such that:*

1. *Contracts maximize the importer's present discounted profits subject to the DICC and the supplier's breakeven constraint.*
2. *Importers choose the origin that maximizes the present discounted value of starting a new relationship.*
3. *The distribution of relationship ages is stationary and satisfies $s_{a+1,odj} = \delta p_a s_{a,odj}$ and $\sum_a s_{a,odj} = 1$.*
4. *Sectoral price indices are consistent with the importers' optimal contracts in ((1) and (2)), labor markets clear (12), and sectoral expenditure is consistent with production and consumption decisions.*

In steady state, destination d chooses to buy variety ω from the origin o that maximizes $V_{0,odj}^b(\omega)$. Consider a shifted marginal cost, $\tilde{c}_{odj} \equiv c_{odj}(\omega)/(\phi_{0,odj}^V)^{1/(\sigma-1)}$. Then, Proposition 1 implies that the chosen origin is simply the origin with the minimum shifted marginal cost, $\arg \min_o \tilde{c}_{odj}(\omega)$. Thus, although origin choice depends on the expected discounted value of a dynamic relationship, scale equivariance makes it equivalent to the familiar minimum-cost sourcing rule of Eaton and Kortum (2002), with delivered marginal cost adjusted by an endogenous origin-specific scalar that summarizes how contracting frictions influence the attractiveness of a given origin.

Proposition 2 (Steady State Trade Shares). *In steady state, the share of destination d 's total supplier*

payments in sector j allocated to origin o is

$$\pi_{odj}^{SS} = \frac{\kappa_{odj}^{-\theta} \bar{\phi}_{odj}^{ext} \bar{\phi}_{odj}^{int}}{\sum_m \kappa_{mdj}^{-\theta} \bar{\phi}_{mdj}^{ext} \bar{\phi}_{mdj}^{int}}, \quad (13)$$

where $\bar{\phi}_{odj}^{ext} \equiv \left(\phi_{0,odj}^V\right)^{\theta/(\sigma-1)-1}$ and $\bar{\phi}_{odj}^{int} \equiv \sum_{a=0}^{\infty} s_{a,odj} \phi_{a,odj}^x$.

Proposition 2 shows that the dynamic contracting frictions enter as aggregate wedges in an otherwise-standard gravity equation. The first wedge, $\bar{\phi}_{odj}^{ext}$, is an extensive margin shifter: d prefers to select origins with lower contracting frictions. The second wedge, $\bar{\phi}_{odj}^{int}$, is an intensive margin shifter: the contracting friction causes d to buy less from a given origin, and the aggregate wedge is simply the age-specific payment wedge, $\{\phi_{a,odj}^x\}_a$, averaged over the stationary distribution of relationship ages. The wedges are functions only of the normalized contracting model and do not depend on any aggregate economic variables, such as τ_{odj} or w_o .

The intuition for why scale equivariance permits tractable aggregation can be easily seen by examining the aggregation explicitly. An equivalent way to express the solution for $x_{a,odj}(\omega)$ in Proposition 1 is to write as shifted relative to the value that would obtain under a frictionless version of the model, $x_{a,odj}(\omega) = \tilde{\phi}_{a,odj}^x x_{odj}^{FB}(\omega)$, with $\tilde{\phi}_{a,odj}^x \propto \phi_{odj}^x$. Substituting this into (10) gives

$$X_{odj} = \sum_a s_{a,odj} \phi_{a,odj}^x \int_0^1 x_{odj}^{FB}(\omega) 1\{\tilde{c}_{odj}(\omega) = \arg \min_m \tilde{c}_{mdj}(\omega)\} d\omega. \quad (14)$$

The intensive margin term inside the integral, $x_{odj}^{FB}(\omega)$, is identical to the Eaton and Kortum (2002) case; the extensive margin term is also identical, with marginal cost simply shifted by an origin-specific scalar. Aggregation therefore follows standard Frechet derivations.

The gravity structure in (13) has two immediate implications. First, the long-run trade elasticity, $-\partial \ln(\pi_{odj}^{SS}/\pi_{ddj}^{SS})/\partial \ln \tau_{odj}$, holding other equilibrium variables fixed, is given by θ . Thus, despite the rich microeconomic dynamics, Frechet-distributed productivity still generates a constant elasticity. Second, if a counterfactual changes trade costs and aggregate economic conditions but leaves the underlying contracting parameters unchanged, then the contracting wedges are identical across steady states. Long-run counterfactuals can therefore be calculated using the standard exact hat algebra approach (Dekle et al., 2008), without requiring knowledge of the levels of the contracting wedges. More broadly, this illustrates that the constant (long-run) elasticity implied by Frechet-distributed heterogeneity in quantitative trade models continues to hold even as one incorporates rich microeconomic features, so long as those microeconomic features satisfy scale equivariance.

2.5 Transition Path

We now turn to the transition path between steady states following a change in trade costs. Suppose that the economy is initially in steady state and, in period $\bar{t}$, it is announced that there is an unanticipated, permanent change in a subset of bilateral trade costs, $\tau_{od} \rightarrow \tau'_{od}$ for some od starting in some period $\bar{t} + s$ for $s \geq 0$. We interpret the shock as an MIT shock, meaning that agents do not foresee it but thereafter have perfect foresight of the path of aggregate economic variables as the economy adjusts to its new steady state.⁷ We show that, although the adjustment involves both the endogenous formation and dissolution of relationships and changes in optimal contracts, the entire transition can be computed in relative changes.

Following the aggregate shock, importers reconsider their choice of origin. While a given importer can in principle leave their incumbent supplier in favor of choosing a new origin immediately in $\bar{t}$ or at any subsequent date, it is without loss to consider only the immediate decision. In particular, backwards induction implies that if the importer plans to exit the relationship in some period t , then it is not possible to satisfy the DICC in period $t - 1$, and the contract unravels. Importers therefore either retain their incumbent supplier until the relationship ends through a unsuitability realization or the good exogenously becomes obsolete.

To simplify analysis, we impose two related restrictions on how the choice of origin occurs. First, once the importer starts a new relationship (which is immediate if they leave their incumbent or following an unsuitable realization if they retain the incumbent), the importer remains in the chosen origin permanently (until the exogenous obsolescence shock occurs). Second, if the importer does not switch, they choose one origin as a designated backup, $b_{dj}(\omega)$. Their outside option is then given by the maximum of a new relationship in the current origin i or a new relationship in the backup, $\max\{V_{0, idj, t}^b(\omega), V_{0, b_{dj}(\omega) dj, t}^b(\omega)\}$. These restrictions continue to allow the outside option in incumbent relationships to vary over time as aggregate economic conditions change, and the decision as to whether to retain the incumbent fully internalizes the future path of these variables. What they prevent is the identity of the outside option changing over time.⁸ This greatly reduces dimensionality as they imply that it is sufficient to track the relative productivity of the incumbent origin and one designated backup, as opposed to the entire vector $\{z_{o, j}(\omega)\}_o$.

History dependence. The first of the two main dynamic mechanisms is that importers enter the transition with existing supplier relationships. In particular, importers have accumulated relationship-specific value, and may prefer to retain their incumbent supplier even when another

⁷Recent work studying dynamic contracts in the presence of aggregate shocks has shown that a coherent and tractable way to incorporate unanticipated shocks into the contract is to allow agents to anticipate the shock occurring with probability p and study the limiting case as $p \rightarrow 0$ (Beraldi, 2025; Donald et al., 2025).

⁸For the counterfactuals involving permanent changes in trade costs that we study in this paper, the identity of the outside option could only change if the changes aggregate endogenous variables are similar in magnitude but different in speed across origins. We view this restriction as a reasonable approximation, as we do not expect general equilibrium movements to be sufficiently volatile that the alternating between two different backups would generate large changes in the outside option.

origin has become more attractive. Consider an importer in a relationship of age a with an incumbent supplier in origin i . The importer retains its incumbent supplier rather than switching to the backup if $V_{a,idj,\bar{t}}^b(\omega) \geq V_{0,b_{dj}(\omega)dj,\bar{t}}^b(\omega)$.

For an importer whose existing origin remains the best choice, retaining the incumbent relationship is always optimal as value is increasing in relationship age. Otherwise, the decision depends on whether the gains from switching outweigh the value accumulated in the existing relationship. We can therefore partition incumbent relationships into three regions: (i) those for which the incumbent origin remains the best choice; (ii) those for which another origin is preferred for a new relationship, but the importer retains its incumbent supplier; and (iii) those for which the importer switches to a new origin and begins a relationship at age zero. Both the second and third region generate gradual adjustment: firms in the third region must build new relationships from scratch with low initial quantities, and firms in the second region will delay switching until the relationship breaks up following an unsuitable realization or exogenous obsolescence (in which case they must subsequently build a new relationship).

Future dependence. The second dynamic mechanism operates through the forward-looking nature of relational contracts. Consider an importer that retains its incumbent supplier. The shock changes not only current trade costs, but also future trade costs, and therefore lowers the future surplus available to sustain repayment. In particular, the DICC for a retained relationship becomes

$$\xi x_{a,idj,t}(\omega) \leq \beta \delta p_a \left(V_{a+1,idj,t+1}^b(\omega) - \max\{V_{0,idj,t+1}^b(\omega), V_{0,b_{dj}(\omega)dj,t+1}^b(\omega)\} \right). \quad (15)$$

The change in trade costs causes $V_{a+1,idj,t+1}^b(\omega)$ to fall relative to the outside option, $V_{0,b_{dj}(\omega)dj,t+1}^b(\omega)$. This tightens the DICC, which reduces the contractible quantity in retained relationships and may cause otherwise profitable retained relationships to break down altogether. This can accelerate the transition relative to a model that contains only the history dependent force, as well as generate anticipatory responses to future changes in economic conditions.

Characterizing the change in outside option. We can characterize the partition into the three regions, as well as how the change in outside option affects the contracts with retained incumbents, as follows. Define the ratio of the fresh-start value in the origin of the incumbent, i , to the fresh-start value in the backup origin, $b_{dj}(\omega)$, as

$$r_{idj,\bar{t}}(\omega) \equiv \frac{V_{0,idj,\bar{t}}^b(\omega)}{V_{0,b_{dj}(\omega)dj,\bar{t}}^b(\omega)} \in (0, 1]. \quad (16)$$

Region (i), where the existing origin remains the best, is the case where $r_{idj,\bar{t}}(\omega) = 1$. There will in general be an age-dependent cutoff $R_{a,ib_{dj}(\omega)dj}$ such that firms with $r_{idj,\bar{t}}(\omega) \in (R_{a,ib_{dj}(\omega)dj}, 1)$ will retain their incumbent supplier (Region (ii)), while firms with $r_{idj,\bar{t}}(\omega) \leq R_{a,ib_{dj}(\omega)dj}$ will switch (Region (iii)). For the retained incumbents, we can use (16) to substitute out the outside option

in (15), which restores scale equivariance up to the ratio r and the identity of the best new origin b . The adjustment of retained contracts therefore differs across varieties in an intuitive way with wedges taking the form $\phi_{a, idj, t}^{x, \text{region } 2}(r, b)$: a lower ratio r means the outside option is more attractive and the contract is further distorted, and the identity b determines how that relative attractiveness varies over time.

Definition 2 (Transition equilibrium). *Given an initial stationary equilibrium, a path of bilateral trade costs $\{\tau_{odj, t}\}_{t \geq \bar{t}}$, and the parameters Θ , a transition equilibrium is a sequence of aggregate variables $\{w_t, P_t, E_t\}_{t \geq \bar{t}}$, optimal contracts and sourcing decisions, and distributions of relationships such that:*

1. *Contracts signed at every date (including those re-signed for retained relationships) maximize the importer's present discounted profits subject to the DICC and supplier's breakeven constraint, taking as given the anticipated future path of aggregate conditions.*
2. *Importers designate a backup origin at the announcement of the shock and choose whether to retain their incumbent or switch to the backup origin immediately to maximize present discounted profit.*
3. *The distribution of relationships evolves following incumbent retention decisions, relationship survival, and exogenous obsolescence.*
4. *At every date, sectoral price indices are consistent with the importers' optimal contracts in ((1) and (2)), labor markets clear (12), and sectoral expenditure is consistent with production and consumption decisions.*

Aggregation. The dynamic mechanisms complicate aggregation relative to the steady state as the three region structure means that we need to integrate not only over the distribution of the minimum (ϕ^V -adjusted) cost but also over the distribution of the ratio of the pre-shock minimum to post-shock minimum, $r_{idj, t}(\omega)$. Nevertheless, two features of the model make this tractable. First, scale equivariance means that, while contracts may depend on (r, b) , they still do not depend on the absolute value of productivity or of the aggregate economic variables. Second, the memorylessness property of the Exponential distribution means that Frechet-distributed productivity permits closed-form aggregation over the ratio $r_{idj, t}(\omega)$ after conditioning on the importer's pre-shock origin choice.⁹

We can decompose bilateral trade at each post-shock date into payments from pre-shock relationships (Regions (i) and (ii)) and payments to relationships formed after the shock (Region (iii) and varieties that became obsolete after the shock). Computing the transition path therefore requires tracking the evolution of the (productivity-weighted) distribution of relationship ages across these

⁹A related use of the statistical relationship between the price in the lowest cost origin and the price in other origins under Frechet draws appears in [Bernard et al. \(2003\)](#), where the second-lowest cost determines pricing under Bertrand competition. Here, the ratio of the pre-shock lowest cost to the post-shock lowest cost determines supplier retention and contract adjustment. The post-shock lowest cost need not have been the pre-shock second-lowest; the general property of memorylessness goes beyond the second order statistic.

groups from the initial period of the shock, $\bar{t}$, through to eventual convergence to the steady state. The following proposition establishes that this can be done entirely in relative changes.

Proposition 3 (Transition Path in Relative Changes). *Denote relative changes in variable x_t as $\hat{x}_t \equiv x_t/x_{t-1}$. Consider an unanticipated shock announced in period $\bar{t}$ in relative changes to trade costs in an economy initially in steady state. For a given path of general equilibrium variables expressed in relative changes, $(\hat{w}_{o,t}, \hat{v}_{dj,t}, \hat{P}_{dj,t})$, values for steady state trade shares, $\{\pi_{odj}\}$, and values for the parameters $(\beta, \delta, \sigma, \theta, \{\gamma\}, \mu, \lambda, \{\xi\})$, the path of aggregate bilateral trade flows, $\{\pi_{odj, \bar{t}+s}\}$ for $s \geq 0$, can be calculated without reference to the levels of bilateral trade shifters $\{\tau_{odj}\}$ or of the general equilibrium variables.*

The explicit characterization of the equations governing the transition path in relative changes are presented in Appendix A. The core conceptual result is that the combination of scale equivariance and the statistical relationship between the pre-shock minimum cost and the post-shock minimum cost implied by Frechet-distributed productivity allows the exact hat algebra approach of Dekle et al. (2008) to be used to calculate the transition path, so the researcher does not need to specify the levels of aggregate economic fundamentals or general equilibrium objects.

Computing the equilibrium transition. Our algorithm to solve for the path of general equilibrium variables in relative changes is similar to that of Caliendo et al. (2019). We assume that the economy reaches the new steady state after a finite number of periods, T . We first conjecture a path of wages, price indices, and demand shifters. Given this path, we solve the contracting problems, determine which incumbent relationships are retained, and iterate forward the distribution of relationship ages. This determines bilateral trade flows, sectoral expenditure, and input demand at every date. We then check if all markets clear and update the conjectured path if they do not.

The transition path features the two dynamic mechanisms described above: history dependence through the distribution of incumbent relationships, and future dependence through anticipated changes in future surplus affecting optimal contracts in the present. In the quantitative analysis, we separate these two mechanisms by comparing the full model with a counterfactual in which the (normalized) wedges induced by the contracting friction are held fixed at their pre-shock values, allowing for violations of the DICC.

3 Empirical Evidence

In this section, we use rich firm-to-firm transaction data from India to measure two aspects of relationship development that play an important role in empirically disciplining the model, and establish two new empirical facts about relationship dynamics that are both pertinent to our model and of interest in their own right.

3.1 Data and Measurement

Our data consists of administrative records of firm-to-firm transactions from India’s Goods and Services Tax (GST) system, obtained from a large Indian state whose population exceeds 100 million. Our main dataset is GSTR-1 filings, which report the universe of domestic firm-to-firm sales by firms registered in the study state, including sales to buyers in other states. The data identify the buyer and seller, invoice date, transaction value, and cover July 2017 through June 2026. We restrict our analysis to transactions involving the 25 largest states by GDP.

We define a relationship as a directed seller-buyer pair and aggregate transactions to the relationship by calendar-quarter level. We define relationship age as the number of prior quarters in which the relationship recorded positive trade.¹⁰ As we cannot measure relationship age for relationships that existed prior to July 2017, we restrict analysis to July 2018 onwards and exclude all pairs that traded in the initial July 2017-June 2018 window, unless otherwise specified.

We complement the GSTR-1 records with data from India’s e-way bill system, for which we observe the universe of bills whether either the seller or buyer is registered in the study state. Introduced in April 2018, firms are required to file an e-way bill for shipments of goods that exceed a value threshold (around USD 500). In addition to recording buyer and seller identifiers and the shipment date, e-way bills record the value, quantity, and product code for each product in the shipment.¹¹ This allows us to separate quantity from price dynamics, and to construct firm-level product histories. However, e-way bills only cover goods and shipments above the threshold. We therefore use the GSTR-1 data for most analysis, and use the e-way bill data when quantities, unit prices, or product codes are required.

3.2 Value Traded and Survival Rate Grow with Relationship Age

The model implies that quantity and value traded, as well as the probability of trading again, grow over the length of a relationship. These patterns have been documented in US international imports (Monarch and Schmidt-Eisenlohr, 2023; Heise, 2024) and among manufacturing firms in Ecuador (Brugués, 2026). We begin our empirical analysis by measuring this value and survival rate growth in our Indian transactions data, which are important for several reasons. First, our novel empirical results relate to how these core facts vary with different dimensions of heterogeneity, such as firm size and local enforcement, so it is useful to establish a baseline. Second, they play an important role in how we estimate the model. Third, relative to prior literature, we explicitly derive our estimating equation from a quantitative model. This makes explicit the appropriate set of fixed effects, clarifies the identification assumption, and allows the coefficients to have a direct structural interpretation.

¹⁰In practice, most relationships (by value and count) trade in every quarter while the relationship is active, so the treatment of gaps has little quantitative effect on our results.

¹¹The granularity of the product codes varies across firms and over time, as reporting thresholds and requirements have changed over time. For most analysis involving unit prices or product histories, we restrict the data to a subset of sellers that consistently report 8-digit HSN product codes.

3.2.1 Growth in Value Traded

Identification To measure the growth in value traded as a function of relationship age, we can take logs of the expression in Proposition 1 to obtain

$$\log x_{a,odj,t}(\omega) = (\sigma - 1) \log z_{o,j}(\omega) + \sigma \log \nu_{dj,t} + (1 - \sigma) \log \kappa_{odj,t} + \log \phi_{a,odj,t}^x. \quad (17)$$

As we do not observe $z_{o,j}(\omega)$, $\nu_{dj,t}$, or $\kappa_{odj,t}$ directly, it is natural to remove them with relationship and market-quarter fixed effects. However, the absolute level of $\phi_{a,odj,t}^x$ is then not identified: for a given odj,t , a scalar that multiplies all $\{\phi_{a,odj,t}^x\}_a$ would be absorbed by the odj,t -level fixed effect. A natural common normalization is to express wedges as relative to a mature benchmark. Recall from Proposition 1 that $\phi_{\infty,odj,t}^x$ is the wedge that obtains if beliefs assign probability one to commercial suitability in period t . Adding and subtracting $\log \phi_{\infty,odj,t}^x$ to (17) gives

$$\log x_{a,odj,t}(\omega) = \underbrace{\log \left(\frac{\phi_{a,odj,t}^x}{\phi_{\infty,odj,t}^x} \right)}_{\text{relative wedge}} + \underbrace{(\sigma - 1) \log z_{o,j}(\omega)}_{\text{relationship FE}} + \underbrace{[\sigma \log \nu_{dj,t} + (1 - \sigma) \log \kappa_{odj,t} + \log \phi_{\infty,odj,t}^x]}_{\text{market-quarter FE}}. \quad (18)$$

Estimating (18) requires us to take a stand on when the mature benchmark is (approximately) attained in practice. We assume that relationships have reached maturity after 12 quarters of trade, so that $\log \left(\phi_{a,odj,t}^x / \phi_{\infty,odj,t}^x \right) = 0$ for $a \geq 12$.¹²

Equation (18) therefore implies the following estimating equation. Let y_{sbt} denote total value sold by seller s to buyer b in calendar quarter t , and Age_{sbt} denote the number of times seller s has sold to buyer b prior to calendar quarter t . Then, we estimate the following equation:

$$\log y_{sbt} = \sum_q \sum_{a=0}^{11} \beta_{a,q} 1\{\text{Age}_{sbt} = a\} 1\{t = q\} + \gamma_{sb} + \gamma_{o(s)d(b)j(s)t} + \varepsilon_{sbt}, \quad (19)$$

where γ_{sb} are relationship fixed effects and $\gamma_{o(s)d(b)j(s)t}$ are fixed effects for origin state by destination state by the seller's HSN4 sector by calendar quarter. Writing $\beta_{a,c}^{odj} \equiv \log \left(\phi_{a,odj,c+a}^x / \phi_{\infty,odj,c+a}^x \right)$, the coefficient $\beta_{a,c}$ in (19) pools $\beta_{a,c}^{odj}$ across market-sectors odj . We aggregate these coefficients across cohorts using cohort relationship shares to obtain an aggregate age profile and report $\exp(\hat{\beta}_a)$. The reported estimates are therefore interpretable structurally as a (log-space) weighted average of the value wedge in a relationship of age a relative to the value that that relationship would trade in the same calendar period if it were mature. Intuitively, (19) is a regression with relationship fixed effects and dummies for ages $\{0, 1, 2, \dots, 10, 11\}$ with ages 12+ serving as the omitted category, and thus identifies the value of trade early in the relationship relative to the value of trade late in the relationship.

We restrict the sample used to estimate (19) to pairs that trade in more than 12 quarters to avoid

¹²While we cannot test this assumption directly, it is re-assuring that the coefficients on ages 10 and 11 are approximately zero.

changes in the sample arising through early exit. The resulting sample in the GSTR-1 data contains 1.74 million relationships. We cluster standard errors at the relationship level. For robustness, we also report estimates of (19) using a version of that includes seller-quarter and buyer-quarter fixed effects to ensure that the growth is not driven by general firm-level growth among either sellers or buyers.

Results We plot the results in Panel (a) of Figure 1. The typical relationship reaches a mature steady state after around eight quarters, and starts at just under half of this mature value. The curve is similar across specifications; if anything, it takes longer to grow when seller-quarter and buyer-quarter fixed effects are included. In Panel (b), we use the e-way bill data to separate value into quantity and price dynamics, we find that the growth in value is entirely driven by growth in quantity. If anything, there is a small price decline in the order of 1-2%. The sign and magnitude of quantity growth is consistent with prior research. The magnitude of price dynamics varies substantially in the literature: [Monarch and Schmidt-Eisenlohr \(2023\)](#) find price declines in US imports in the order of 30%, while both [Brugués \(2026\)](#) and [Fontaine et al. \(2023\)](#) find declines in the order of 3%-4% among manufacturers in Ecuador and importers in France, respectively. Our estimates are consistent with the latter.

3.2.2 Growth in Survival Rate

To estimate the survival rate as a function of relationship age, we define an indicator D_{sbt} that is 1 if relationship sb ever trades again after period t . We then estimate the following regression:

$$D_{sbt} = \sum_{a=1}^{11} \beta_a 1\{\text{Age}_{sbt} = a\} + \gamma_{d(b)t} + \gamma_{j(s)t} + \varepsilon_{sbt}. \quad (20)$$

Intuitively, this regression estimates the probability that a relationship that survives to age a ever trades again relative to a relationship of age 12+. As we are interested here in absolute rather than relative levels, we used the fitted coefficients to predict the absolute probability of trading again. We exclude observations in the last four quarters of the data to avoid artificially declaring relationships as terminated simply because our data does not cover subsequent trading periods.


We plot the estimated survival rates in Figure 2. The probability that any given pair that has only traded once ever trades again is around 0.5. This rises monotonically and eventually peaks at a long-run survival rate of around 95%. The magnitude is in line with prior research.

3.3 New Empirical Patterns

3.3.1 Near-invariance of the Value Growth Curve to Firm Heterogeneity


In this subsection, we examine the extent to which the value growth curve presented in Panel (a) of Figure 1 varies with several important margins of firm heterogeneity. This is important

Figure 1: Relationship Dynamics: Value Growth Curve



Note: Panel (a) shows value growth over the course of a relationship, estimated using the specification in (19). Panel (b) decomposes value growth in the main specification into quantity and price dynamics.

Figure 2: Relationship Dynamics: Survival Curve



Note: This figure shows the estimated probability of a relationship trading again, conditional on reaching a given age, estimated using (20).

for two reasons. First, it helps to assess the prediction of scale equivariance from the contracting model, which is central to tractable aggregation. This property predicts that, while firms may have heterogeneous productivities that affect their absolute scale of trade, the proportional buildup should be constant across firms. Second, understanding how this curve varies is of more general interest, as different theories as to why this curve exists predict different patterns of heterogeneity. For example, learning about market conditions or demand for a product would suggest that there should be less build-up when the firm is older or has prior experience with the product. The fact that we observe both the universe of within-state transactions, where we have considerable information about both parties, and the universe of sales to other states, over a long horizon allows us to construct several measures of heterogeneity.

To study this, we divide firms into groups based on size, age, location experience, and product experience. We describe the construction and any sample restrictions of each of these four margins of heterogeneity below. In all cases, we augment the value growth specification in (19) to include interaction terms for the relevant margin of heterogeneity, as follows:

$$\begin{aligned} \log y_{sbt} = & \gamma_{sb} + \gamma_{o(s)d(b)j(s)t} + \sum_q \sum_{a=0}^{11} \beta_{a,q} 1\{\text{Age}_{sbt} = a\} 1\{t = q\} \\ & + \sum_{a=0}^{11} \eta'_a \text{Het}_{sb} 1\{\text{Age}_{sbt} = a\} + \varepsilon_{sbt}, \end{aligned} \quad (21)$$

where Het_{odt} captures the relevant interaction variable. As described below, in some cases we include additional controls or fixed effects.

Firm Size To ensure we can accurately measure size of both the buyer and seller, we restrict the sample to buyer-seller pairs in the study state. Within each HSN-4 sector, we divide firms into four quartiles based on their total sales in July 2017 through June 2018, with firms with zero sales in this window excluded. As the analysis already drops all relationships with any positive trade from the value growth regression, there is no mechanical link between total sales in this window and trade flows in the regression.

Firm Age We measure firm age using the earliest date at which the firm registered for value-added taxes in India. For firms that first registered after the introduction of the Goods and Services Tax (GST) in July 2017, we observe their GST registration date. For firms already operating when GST was introduced, we instead observe their registration date under the previous VAT system. As the VAT system was introduced in 2005, the registration date for firms that predate it is left-censored at 2005. We divide firms into bins based on their registration date and include heterogeneous slopes by HSN-4 sector to control for age differences correlated with sector. As we only observe this registration information for firms in the study state, we again restrict the sample to buyer-seller pairs where both parties are in the study state.

Location Experience A large literature studies how values or quantities evolve as a function of time since a firm first exported to or imported from a new market, with several prominent theories of these dynamics emphasizing learning about market-specific conditions such as demand or quality (e.g., [Berman et al. \(2019\)](#); [Fitzgerald et al. \(2024\)](#)). As some location-specific learning mechanisms would predict that relationships may start closer to maturity if the firm has location-specific experience, we consider heterogeneity by whether a firm has previously sold to a given state (for sellers) and whether a firm has previously bought from a given state (for buyers). We restrict the sample to interstate sales for sellers in our study state, and interstate purchases for buyers in our study state. As we only observe interstate purchases in the e-way bill data, for comparability we restrict both analyses to the e-way bill data.

Product Experience Another literature posits that firm export and import dynamics reflect the firm learning things about the product itself that are transferable across trading partners but are not market-specific, such as whether a purchased input is compatible or whether there is demand for an output (e.g., [Eaton et al. \(2025\)](#)). This predicts that growth should be slower for firms that already have experience with the product. We use the e-way bill data, where we observe product codes, and restrict to a sub-sample of sellers that consistently report 8-digit HSN codes. For every relationship, we can construct an indicator that is one if the firm has sold (for sellers) or bought (for buyers) this product in our data. We again restrict to trade within the study state to ensure we always observe the full sales and purchase history.


We plot the results in Figure 3. The main finding is that the value growth curve is remarkably consistent across these important margins of firm heterogeneity. While there are occasional differences on some variables at some ages, these are generally small in magnitude—on the whole, the evidence shows that relationship growth is a pervasive phenomenon across the economy, rather than something confined to a particular group of firms. We interpret this as empirical support for the separability between the scale of trade and relationship dynamics.

3.3.2 Relationship Value Growth and Contract Enforcement


While we have shown that the value growth curve is remarkably consistent across firm characteristics and experience, the model predicts that it should vary with the strength of contract enforcement. In this section, we therefore examine empirically whether early trade is closer to mature trade when enforcement is stronger. This provides identifying variation for the contracting parameter ξ , but also provides direct reduced form evidence as to the mechanisms that explain the value growth curve.

To do this, we exploit the fact that the local contracting environment varies substantially across Indian states. In particular, we follow the identification strategy that [Boehm and Oberfield \(2020\)](#) (BO) use to study the effect of contract enforcement on the decision of firms' outsourcing decisions. BO measure court congestion using the average age of pending civil cases in each High Court in


Figure 3: (Lack of) Heterogeneity in Relationship Dynamics




(a) Seller Size




(b) Buyer Size




(c) Seller Age




(d) Buyer Age




(e) Seller Destination Experience



(f) Buyer Origin Experience



(g) Seller Product Experience



(h) Buyer Product Experience

Note: Each panel shows the evolution of value traded over the course of a buyer–seller relationship, separately by seller or buyer characteristics. Panels (a) and (b) examine firm size; panels (c) and (d) examine firm age; panels (e) and (f) examine experience selling to the destination state (for sellers) and experience buying from the origin state (for buyers); and panels (g) and (h) examine experience selling or buying the HSN8-level product.

2016, using an instrumental variables strategy that exploits the historical origins of India's High Courts. Throughout India's history, states have frequently been divided or reorganized, often in response to political movements. These reorganizations were often accompanied by the establishment of new High Courts. Newly established courts typically began with little inherited backlog, while older courts continued to accumulate cases.

The exclusion restriction requires that the timing of these historical state reorganizations is not correlated with the current growth rate of buyer-seller relationships through channels other than court congestion. To partly address concerns arising from the possibility that newer states are systematically different, we allow the relationship age profile to vary with GDP per capita, population density, distance from the study state, and population. We also allow for heterogeneous slopes by the seller's 4-digit HSN sector. BO investigate the validity of the instrument in detail, and show that their results are robust to controlling for development, corruption, and ethnolinguistic fragmentation.

We examine how the relationship growth curve varies with court congestion by interacting the court variables with the relationship age indicators in (19). We restrict the sample to interstate relationships to avoid attributing differences in own-state versus interstate trade to differences in enforcement. Defining $Z_{d(b)}$ as the High Court age in the destination state (where the buyer is located), the reduced form is the following:


$$\begin{aligned} \log y_{sbt} = & \gamma_{sb} + \gamma_{j(s)t} + \sum_q \sum_{a=0}^{11} \beta_{a,q} 1\{\text{Age}_{sbt} = a\} 1\{t = q\} \\ & + \sum_{a=0}^{11} (\eta_a Z_{d(b)} + \rho'_a X_{o(s),d(b)} + \gamma_{j(s),a}) 1\{\text{Age}_{sbt} = a\} + \varepsilon_{sbt}, \end{aligned} \quad (22)$$

where $X_{o(s),d(b)}$ is a vector of state- and state-pair level development covariates, and $\gamma_{j(s),a}$ are fixed effects for the seller's 4-digit HSN sector. The second-stage specification is analogous, with predicted backlog replacing $Z_{d(b)}$. The coefficients of interest are $\{\eta_a\}$, which capture how the shape of the value growth curve varies with the court variables. Formally, we use 2SLS, clustering standard errors at the level of the High Court.

We plot the results in Figure 4. Panel (a) shows the interaction coefficients, and Panel (b) shows the predicted value growth curve at different levels of enforcement. The interaction coefficients are significant at either the 5% or 10% level and of economically important magnitude.¹³ In Panel (b), comparing the predicted growth curve among courts towards the top of the enforcement distribution (two standard deviations above the mean) to the average growth curve shows that stronger courts meaningfully flatten the curve, reaching maturity after roughly one year as compared to roughly two years for the average curve.

¹³The differences in the curves in Panel (b) are not simply $\pm$ the interaction coefficients in Panel (a). This is because the interaction coefficients are in log space, while the curves in Panel (b) are in level space.

Figure 4: Relationship Dynamics and Destination Contract Enforcement



Note: Panel (a) shows the interaction coefficients. Panel (b) shows the predicted value growth curve at different levels of contract enforcement.

If the pattern in Figure 4 reflects the causal effect of improved contract enforcement, a natural test is to examine whether the effect is larger in sectors where goods are more relationship-specific and therefore more subject to hold-up problems. To examine this, we estimate a version of (22) where we additionally interact the court heterogeneity variables with the measure of relationship-specificity from Rauch (1999).¹⁴ We define the relationship's sector as the HSN-4 code with the highest value in the first quarter that the pair trades. We plot the results in Figure 5. The coefficients are positive and significant, implying that stronger enforcement does indeed have a larger effect on the value growth curve in sectors with goods that are more relationship-specific. The magnitudes on the triple interaction of around 0.02-0.04 are sizable compared to the magnitudes of 0.05-0.1 from the specification in Figure 4 that pools sectors. The interaction coefficients are less downward-sloping than in Figure 4, although the standard errors are relatively wide.

4 Estimation


We now estimate the model using the Indian firm-to-firm transaction data in addition to new data on aggregate interstate trade. We calibrate standard parameters governing discounting, the demand elasticity, and the trade elasticity, and estimate the parameters governing relationship buildup and the contracting friction using the empirical results presented in Section 3.

4.1 Quantitative Economy and Calibrated Parameters

The economy comprises $D = 20$ locations: the 19 largest Indian states by GDP and a residual location that aggregates the remaining states and union territories. We work with $S = 15$ sectors using

¹⁴Rauch (1999) classifies goods into relationship-specific or not based on whether they are sold on organized exchanges. Boehm and Oberfield (2020) also use this measure to study whether states with weaker enforcement see more distortion in outsourcing decisions specifically within sectors that are more relationship-specific.

Figure 5: Triple Interaction Coefficients: Age $\times$ Enforcement $\times$ Relationship-Specificity



Note: This figure shows the interaction coefficients from estimating the value growth curve specification with interactions covering destination-level court congestion (instrumented with High Court Age) and the relationship-specificity of the goods from [Rauch \(1999\)](#).

the standard definition in India KLEMS. For the bilateral trade flow matrix, we use an additional proprietary dataset covering aggregate bilateral flows between at the state by HSN-4 product by month level, which we aggregate to the final quarter of the fiscal year 2022-23 and take to be our baseline period.

We construct the cost share parameters governing input-output linkages, $\{\gamma_{jk}\}$, using the Supply Use Tables published by the Government of India’s National Statistics Office. We construct household expenditure shares, $\{\alpha_{d,j}\}$, using India’s Household Consumption Expenditure Survey.

To match our empirical analysis in Section 3, we set a model period to correspond to a calendar quarter. Following [Caliendo et al. \(2019\)](#), we set the quarterly discount factor β to 0.99. We set the CES demand elasticity σ to 3, following standard practice in the macroeconomics literature (e.g., [Hsieh and Klenow \(2009\)](#)), and the Frechet shape parameter θ to 4 following [Barnwal et al. \(2026\)](#). We set the exogenous non-obsolence continuation probability δ equal to the survival probability of mature relationships of 0.95 from the survival curve estimated in Figure 2.

4.2 Estimating Learning and Enforcement Parameters

We jointly estimate the parameters governing commercial suitability and enforcement to match the 12 age coefficients in the value growth curve in Figure 1 and the 12 interaction coefficients in the court heterogeneity regression in Figure 4. For simplicity, we assume that the parameters (μ, λ) that govern learning about commercial suitability are constant across origins, destinations,

and sectors. As our model emphasizes a hold-up problem on the buyer side, and our empirical evidence identifies heterogeneity in the value growth curve by enforcement in the destination state, we assume that the contracting friction parameter, ξ , varies at the destination level. As we do not know the exact relationship between court congestion and $\{\xi_d\}$, we impose a two parameter logistic relationship between the two:

$$\log \left(\frac{\xi_d}{1 - \xi_d} \right) = \eta_0 + \eta_1 \text{CourtCongestion}_d, \quad (23)$$

where $\eta_1 > 0$ implies that more congested courts increase the share of value that the buyer can gain by holding up the seller. Reducing the estimation of $\{\xi_d\}$ to two parameters, rather than estimating ξ_d directly for every state, reflects the fact that the interaction coefficients use cross-sectional variation across destinations to identify the relationship between the growth curve and courts.


To understand the connection between the model and the empirical moments, consider payments at two ages a and a' in the same relationship within a stationary environment. Proposition 1 and Equation (8) together imply

$$\frac{x_{a,odj}}{x_{a',odj}} = \frac{\phi_{a,odj}^x}{\phi_{a',odj}^x}, \quad x_{a,odj} = \min \left\{ \underbrace{\left(\frac{\sigma - 1}{\sigma} \right)^\sigma (p_a \nu_{dj})^\sigma c_{odj}^{1-\sigma}}_{\text{learning-implied } x_a}, \underbrace{\frac{\beta \delta p_a (V_{a+1}^b - V_0^b)}{\xi_d}}_{x_a \text{ if DICC binds at } a} \right\}. \quad (24)$$

The equation on the left is a consequence of scale equivariance: the producer-specific productivity $z_{o,j}(\omega)$ and the demand and cost shifters ν_{dj} and c_{odj} cancel. Assuming that any subsequent effects of learning from age 12 onwards are numerically small, taking a' to be the 12+ bin used in Figure 1 shows that the relationship growth curve identifies the wedges, $\phi_{a,odj}^x$, at every age a relative to the long-run wedge, $\phi_{\infty,odj}^x$. The equation on the right maps wedges to parameters. Recall from (8) that quantity (and thus value) is the minimum of the optimal order size given current beliefs $p_a = 1 - \mu_a + \mu_a \lambda$ and the largest order size consistent with the DICC at age a . As the learning about commercial suitability is constant across destinations, variation in local contract enforcement, ξ_d , generates variation in the shape of the growth curve by changing both the value of the DICC term and the ages at which it is active. Intuitively, when courts are weaker, the DICC is more binding, meaning that $x_{a,odj}$ is lower than the learning model alone would imply. Thus, the extent to which the growth curve varies with court enforcement identifies (η_0, η_1) .

Write the vector of parameters to be estimated as $\psi \equiv (\mu, \lambda, \eta_0, \eta_1)$. For each candidate ψ , we solve the contracting problem to construct the model-implied counterparts to the empirical growth curve and court heterogeneity coefficients, denoted $m_a^V(\psi)$ and $m_a^C(\psi)$. We then use the following

Figure 6: Targeted Moments: Model vs Data



minimum distance estimator:

$$\hat{\psi} = \arg \min_{\psi} \left\{ \sum_{a=1}^{11} w_a^V [m_a^V(\psi) - \hat{m}_a^V]^2 + \sum_{a=1}^{11} w_a^C [m_a^C(\psi) - \hat{m}_a^C]^2 \right\}, \quad (25)$$

where $\hat{m}_a^V$ and $\hat{m}_a^C$ are the empirical moments, and w_a^V and w_a^C are within-block weights that are proportional to the inverse of the empirical variance of the coefficients, normalized such that the total within-block weight is 1. The normalization gives the two blocks equal total weight, while placing greater weight on more precisely estimated moments within each block.

We obtain the following estimates: $(\hat{\mu}, \hat{\lambda}, \hat{\eta}_0, \hat{\eta}_1) = (0.31, 0.60, -1.30, 0.29)$. We plot the fit on the targeted moments in Figure 6.

5 Quantitative Application

We use the estimated model to study adjustment to a permanent reduction in interstate trade costs. Lowering the costs of trading across space has been an important policy objective in India (as well as in other countries and regions) for decades, with major national policies including the flagship 2017 Goods and Services Tax (GST) reform that harmonized value-added taxes across states, numerous high value national highway infrastructure projects, and a network of dedicated freight corridors that exclusively serve freight trains.

5.1 The Goods and Services Tax Reform

To simulate a reduction in trade costs of realistic magnitude, we consider a policy experiment calibrated to the 2017 GST reform. [Barnwal et al. \(2026\)](#) use high-frequency GPS data on truck movements to show that the GST reform substantially reduced delays at state-border checkpoints, and use the road network to translate those changes into reductions in iceberg trade costs. We

follow this method and their preferred calibration.¹⁵ This provides a matrix of reductions in the transportation component of bilateral trade costs that has a mean value of τ'_{odj}/τ_{odj} of 0.9722 (so an average reduction of around 3%). Own-state trade costs, τ_{odj} for $o = d$, are unchanged.

The shock is unanticipated and implemented immediately. After its announcement, firms correctly anticipate the subsequent path of aggregate economic conditions. We initialize the economy at the stationary distribution of relationships implied by the estimated model and solve for the new steady state and the transition path between the two equilibria.

5.2 Isolating the Role of Forward-Looking Contracts

The model features two dynamic adjustment mechanisms. First, the presence of incumbent relationships and the need to build new relationships upon switching can generate slow adjustment. Second, the fact that relational contracts depend on future surplus means that incumbent relationships that would otherwise remain profitable can reduce quantity or collapse altogether if the improvement in the outside option, which lowers the present and future surplus, tightens the DICC sufficiently. To separate these, we consider a version of the same counterfactual, but where we hold the normalized relationship wedges fixed at their pre-shock values: $\phi_{a,odj,t}^x = \phi_{a,odj}^{x,SS}$. This corresponds to a view in which the wedges need not be generated endogenously by contracting, but are simply exogenous technological features that govern relationship dynamics in the economy. Scale equivariance continues to hold, so quantity and value traded continue to vary as demand and cost conditions change, and new relationships still begin at age zero and mature over time. Thus, in both the full model with endogenous ϕ and the version with fixed ϕ , the transition features relationship buildup, endogenous supplier retention, and forward-looking switching decisions. What the comparison isolates is the endogenous response of ϕ to changes in future surplus, which includes the possibility that a given incumbent relationship may have been incentive-compatible in the steady state but not in the new regime.


5.3 Results

To summarize the adjustment of bilateral trade, and to relate to recent empirical evidence showing that the short-run trade elasticity is substantially lower than the long-run trade elasticity (Boehm et al., 2023), we calculate horizon-specific trade elasticities under the model. At each horizon h following the shock, we estimate the following gravity regression using model-generated flows:

$$\log \left(\frac{X_{odj,\bar{t}+h}}{X_{odj}^{SS}} \right) = \gamma_{ojh} + \gamma_{djh} - \varepsilon_h \log \hat{\tau}_{odj} + u_{odjh}, \quad (26)$$

¹⁵In particular, we assume that iceberg trade costs have a multiplicative component given by $\tau_{od} = 1 + \gamma T_{od}^\rho$, where T_{od} is travel time. Barnwal et al. (2026) set $\rho = 0.8$ calibrate γ so that the median pre-reform trade cost equals 2.5. Then, we use the 2016 road network to construct travel times at baseline, and their estimated reduction in time spent at border checkpoints of 32 minutes (from a base of 52 minutes) to construct travels times after the reform.

Figure 7: Counterfactual Trade Adjustment



where γ_{ojh} and γ_{djh} are origin-sector-horizon and destination-sector-horizon fixed effects, respectively. As shown in Proposition 2, the long-run elasticity is the Frechet shape parameter θ , which in our case is 4, so in the long run we will have $\varepsilon_h \rightarrow 4$. As discussed in Section 2, the two dynamic mechanisms primarily affect the extensive margin and therefore the short-run elasticity will generally be no lower than the intensive margin, which is $\sigma - 1 = 2$. The question is therefore the extent to which elasticity in short horizons is closer to 2 than to 4.

We plot the results in Panel (a) of Figure 7. In line with the evidence in Boehm et al. (2023), under both the full and the fixed-wedge model the short-run elasticity is substantially lower than the long-run elasticity. In the full model with endogenous wedges, the immediate elasticity is 2.85, while in the model with fixed wedges it is 2.52. After five years the respective elasticities are 3.67 and 3.48, still short of the long-run value of 4. The elasticity reaches 90% of its long-run value after 15 quarters in the full model and after 30 quarters in the fixed-wedge model, so freezing the contract wedges doubles the length of the adjustment.

In Panel (b) of Figure 7 we display the share of firms that, if starting from scratch would choose a different origin, but choose to retain their incumbent. There are two relevant observations. First, under both models these retention shares are large: 67% in the full model and 88% in the fixed-wedge model. This highlights that the history dependence of relationships is quantitatively important for the extensive margin of sourcing: many buyers retain their incumbent supplier because of (endogenously) accumulated relationship capital. Second, the retention share is substantially lower in the full model. When outside options improve, the surplus in incumbent relationships falls, and a large mass of otherwise profitable relationships break down, tempering the history-dependent incumbent retention channel.

6 Conclusion

In this paper, we study how the dynamics of buyer-seller relationships affect aggregate adjustment to changes in trade costs. We develop a dynamic quantitative trade model in which relationships grow gradually through learning, generating history dependence, while imperfect contract enforcement means that the scale of current trade depends on future surplus, generating future dependence. We show that these mechanisms can be incorporated tractably into a multi-location, multi-sector general equilibrium model through a property that we call scale equivariance.

Using rich firm-to-firm transaction data from India, we document that the proportional buildup of trade is remarkably similar across firm characteristics and experience, but varies with the strength of contract enforcement. We use these empirical results to estimate the model and study adjustment to a reduction in interstate trade costs. We find that relationships substantially slow trade adjustment, as importers retain incumbent suppliers and new relationships take time to develop. Offsetting this, the changes in surplus implied by improved outside options weakens incentives and causes some otherwise profitable incumbent relationships to dissolve. Together, these results suggest that both the history-dependent and future-dependent components of relationship dynamics are important for predicting and understanding the dynamics of trade adjustment.

References

- P. Antràs. Firms, contracts, and trade structure. *The Quarterly Journal of Economics*, 118(4):1375–1418, 2003.
- P. Antras and C. F. Foley. Poultry in motion: a study of international trade finance practices. *Journal of Political Economy*, 123(4):853–901, 2015.
- L. Araujo, G. Mion, and E. Ornelas. Institutions and export dynamics. *Journal of International Economics*, 98:2–20, 2016.
- P. Barnwal, J. I. Dingel, D. Iurchenko, P. Krishna, and E. Van Leemput. Internal trade barriers in india. 2026.
- F. Beraldi. Banking relationships and loan pricing disconnect. 2025.
- N. Berman, V. Rebeyrol, and V. Vicard. Demand learning and firm dynamics: evidence from exporters. *Review of Economics and Statistics*, 101(1):91–106, 2019.
- A. B. Bernard, J. Eaton, J. B. Jensen, and S. Kortum. Plants and productivity in international trade. *American economic review*, 93(4):1268–1290, 2003.
- A. Blouin and R. Macchiavello. Strategic default in the international coffee market. *The Quarterly Journal of Economics*, 134(2):895–951, 2019.

- C. E. Boehm, A. A. Levchenko, and N. Pandalai-Nayar. The long and short (run) of trade elasticities. *American Economic Review*, 113(4):861–905, 2023.
- J. Boehm and E. Oberfield. Misallocation in the market for inputs: Enforcement and the organization of production. *The Quarterly Journal of Economics*, 135(4):2007–2058, 2020.
- J. Boehm, R. South, E. Oberfield, and M. Waseem. The network origins of firm dynamics: Contracting frictions and dynamism with long-term relationships. *NBER Working Paper*, 2024.
- F. Brugués. Take the goods and run: Contracting frictions and market power in supply chains. *American Economic Review*, 116(2):582–626, February 2026.
- L. Caliendo and F. Parro. Estimates of the trade and welfare effects of nafta. *The Review of Economic Studies*, 82(1):1–44, 2015.
- L. Caliendo, M. Dvorkin, and F. Parro. Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica*, 87(3):741–835, 2019.
- J. Chen. Sourcing frictions meet inventories a dynamic ricardian framework for the impact of trade shocks. 2025.
- J. Chen, C. C. Góes, M.-A. Muendler, and F. Trottner. Dynamic adjustment to trade shocks. Technical report, National Bureau of Economic Research, 2026.
- D. Chor and L. Ma. The aggregate welfare and trade implications of contracting frictions in global sourcing. Technical report, National Bureau of Economic Research, 2025.
- S. Das, M. J. Roberts, and J. R. Tybout. Market entry costs, producer heterogeneity, and export dynamics. *Econometrica*, 75(3):837–873, 2007.
- R. Dekle, J. Eaton, and S. Kortum. Global rebalancing with gravity: Measuring the burden of adjustment. *IMF Staff Papers*, 55(3):511–540, 2008.
- E. Donald, M. Fukui, and Y. Miyauchi. Optimal dynamic spatial policy. Technical report, National Bureau of Economic Research, 2025.
- J. Eaton and S. Kortum. Technology, geography, and trade. *Econometrica*, 70(5):1741–1779, 2002.
- J. Eaton, M. Eslava, D. Jenkins, C. Krizan, and J. Tybout. A search and learning model of export dynamics. *Journal of International Economics*, page 104155, 2025.
- D. Fitzgerald, S. Haller, and Y. Yedid-Levi. How exporters grow. *Review of Economic Studies*, 91(4):2276–2306, 2024.
- F. Fontaine, J. Martin, I. Mejean, et al. Frictions and adjustments in firm-to-firm trade. Technical report, Centre for Economic Policy Research, 2023.

- G. M. Grossman, E. Helpman, and S. J. Redding. When tariffs disrupt global supply chains. *American Economic Review*, 114(4):988–1029, 2024.
- A. Harris and T. M. A. Nguyen. Long-term relationships in the us truckload freight industry. *American Economic Journal: Microeconomics*, 17(1):308–353, 2025.
- S. Heise. Firm-to-firm relationships and the pass-through of shocks: Theory and evidence. *Review of Economics and Statistics*, pages 1–45, 2024.
- C.-T. Hsieh and P. J. Klenow. Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448, 2009.
- B. Kleinman, E. Liu, and S. J. Redding. Dynamic spatial general equilibrium. *Econometrica*, 91(2):385–424, 2023.
- B. Kleinman, E. Liu, and S. J. Redding. Dynamic effects of trade. Technical report, National Bureau of Economic Research, 2026.
- D. Kohn, E. Luttini, M. Szkup, and S. Zhang. International trade finance and learning dynamics. *Journal of International Economics*, 162:104286, 2026.
- R. Macchiavello and A. Morjaria. The value of relationships: evidence from a supply shock to kenyan rose exports. *American Economic Review*, 105(9):2911–2945, 2015.
- R. Macchiavello and A. Morjaria. Competition and relational contracts in the rwanda coffee chain. *The Quarterly Journal of Economics*, 136(2):1089–1143, 2021.
- A. Marcet and R. Marimon. Recursive contracts. *Econometrica*, 87(5):1589–1631, 2019.
- M. J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.
- R. Monarch and T. Schmidt-Eisenlohr. Longevity and the value of trade relationships. *Journal of International Economics*, 145:103842, 2023.
- N. Nunn. Relationship-specificity, incomplete contracts, and the pattern of trade. *The quarterly journal of economics*, 122(2):569–600, 2007.
- J. E. Rauch. Networks versus markets in international trade. *Journal of international Economics*, 48(1):7–35, 1999.
- E. Wiles and D. Houeix. Relational frictions along the supply chain: Evidence from senegalese traders. *Available at SSRN 5397962*, 2026.

A Theoretical Appendix

A.1 Proof of Proposition 1

Proposition 1 (Scale equivariance with same-origin rematching). *Consider the contracting problem in a buyer-seller match. Suppose that the buyer's outside option is given by re-matching in the same origin, and that there exists a solution to the contracting problem consistent with this. Then, its solution can be expressed in normalized units that depend only on age, the structural parameters, and the future relative paths $\{\nu_t/\nu_{t_0}, \kappa_t/\kappa_{t_0}\}_{t \geq t_0}$. In particular, there exist scalars $\phi_{a,odj,t}^q, \phi_{a,odj,t}^x, \phi_{a,odj,t}^V$ such that*

$$\begin{aligned} q_{a,odj,t}(\omega) &= z_{o,j}(\omega)^\sigma \nu_{djt}^\sigma \kappa_{odjt}^{-\sigma} \phi_{a,odj,t}^q \\ x_{a,odj,t}(\omega) &= z_{o,j}(\omega)^{\sigma-1} \nu_{djt}^\sigma \kappa_{odjt}^{1-\sigma} \phi_{a,odj,t}^x \\ V_{a,odj,t}(\omega) &= z_{o,j}(\omega)^{\sigma-1} \nu_{djt}^\sigma \kappa_{odjt}^{1-\sigma} \phi_{a,odj,t}^V. \end{aligned}$$

Proof. Consider a given importer-supplier match and suppress all aggregate indices. Let the sequence of outside options be given by $O \equiv \{O_t\}_{t \geq t_0}$. Let productivity be z , so marginal cost is $c_t = \kappa_t/z$. The contract specifies quantities $q_{a,t}$ and transfers $x_{a,t}$. The buyer's continuation value is $V_{a,t}$.

Consider multiplying productivity by $\lambda > 0$, the entire future path of cost shifters by $K > 0$, and the entire future path of demand shifters by $N > 0$. Thus, $z' = \lambda z$, $\kappa'_t = K \kappa_t$, and $\nu'_t = N \nu_t$ for every $t \geq t_0$. Suppose that the outside option is transformed according to $O'_t = \lambda^{\sigma-1} N^\sigma K^{1-\sigma} O_t$ for all t .

Consider transforming every quantity, payment, and continuation value according to

$$q'_{a,t} = \lambda^\sigma N^\sigma K^{-\sigma} q_{a,t}, \quad x'_{a,t} = \lambda^{\sigma-1} N^\sigma K^{1-\sigma} x_{a,t}, \quad V'_{a,t} = \lambda^{\sigma-1} N^\sigma K^{1-\sigma} V_{a,t}. \quad (27)$$

Observe that both the buyer's and the seller's flow payoffs are simply scaled by $S \equiv \lambda^{\sigma-1} N^\sigma K^{1-\sigma}$:

$$\pi_{a,t}^{b'} = \nu'_t (q'_{a,t})^{(\sigma-1)/\sigma} - x'_{a,t} = N \nu_t (\lambda^\sigma N^\sigma K^{-\sigma} q_{a,t})^{(\sigma-1)/\sigma} - \lambda^{\sigma-1} N^\sigma K^{1-\sigma} x_{a,t} = S \pi_{a,t}^b \quad (28)$$

$$\pi_{a,t}^{s'} = x'_{a,t} - \kappa'_t / (z') q'_{a,t} = S x_{a,t} - S \kappa_t / z q_{a,t} = S \pi_{a,t}^s. \quad (29)$$

The learning process is unchanged as success probabilities and beliefs do not depend on z . Equation (29) implies that the seller's breakeven constraint is unchanged. The buyer's incentive constraint is also unchanged:

$$\xi x'_{a,t} \leq \beta \delta p_a \left(V'_{a+1,t+1} - O'_{t+1} \right) \iff \xi S x_{a,t} \leq \beta \delta p_a \left(S V_{a+1,t+1}^b - S O_{t+1} \right). \quad (30)$$

Therefore, any contract that is feasible and incentive compatible with (z, ν, κ) under outside option O has a one-to-one transformed counterpart that is feasible and incentive compatible with

productivity (z', ν', κ') under outside option O' . Moreover, the transformation preserves rankings across contracts. Write the *normalized* optimal contract as the optimal contract in the special case of $z = \nu_{t_0} = \kappa_{t_0} = 1$ as $\{(\tilde{q}_{a,t}, \tilde{x}_{a,t})\}$. Then, for any $(z, \nu_{t_0}, \kappa_{t_0})$ the optimal contract policies and corresponding continuation value are given by

$$\begin{aligned} q_{a,t}(z, \nu_{t_0}, \kappa_{t_0}; O) &= z^\sigma \nu_{t_0}^\sigma \kappa_{t_0}^{-\sigma} \tilde{q}_{a,t}(O) \\ x_{a,t}(z, \nu_{t_0}, \kappa_{t_0}; O) &= z^{\sigma-1} \nu_{t_0}^\sigma \kappa_{t_0}^{1-\sigma} \tilde{x}_{a,t}(O) \\ V_{a,t}^b(z, \nu_{t_0}, \kappa_{t_0}; O) &= z^{\sigma-1} \nu_{t_0}^\sigma \kappa_{t_0}^{1-\sigma} \tilde{V}_{a,t}(O). \end{aligned}$$

The contracts stated above are for an exogenous outside option that satisfies transformation $O' = SO$. It remains to be shown that, if the outside option is defined as rematching, that this scales as required. Define $F_t(O) \equiv V_{0,t}^b(z, \nu_{t_0}, \kappa_{t_0}; O)$ as the value of beginning new relationship implied by the solved contract that takes O as given. As stated in the Proposition, the original (untransformed) contract satisfies $O_t = F_t(O)$ for all $t \geq t_0$. Observe that F_t scales with S , which implies $F_t'(O') = SF_t(O) = SO_t = O'_t$. Thus, the outside option scales in the manner conjectured.

To obtain solutions expressed in terms of contemporaneous values of (ν_t, κ_t) , define

$$\phi_{a,t}^q = \left(\frac{\nu_{t_0}}{\nu_t}\right)^\sigma \left(\frac{\kappa_{t_0}}{\kappa_t}\right)^{-\sigma} \tilde{q}_{a,t}, \quad \phi_{a,t}^x = \left(\frac{\nu_{t_0}}{\nu_t}\right)^\sigma \left(\frac{\kappa_{t_0}}{\kappa_t}\right)^{1-\sigma} \tilde{x}_{a,t}, \quad \phi_{a,t}^V = \left(\frac{\nu_{t_0}}{\nu_t}\right)^\sigma \left(\frac{\kappa_{t_0}}{\kappa_t}\right)^{1-\sigma} \tilde{V}_{a,t}.$$

Then, we obtain

$$\begin{aligned} q_{a,t}(\omega) &= z(\omega)^\sigma \nu_t^\sigma \kappa_t^{-\sigma} \phi_{a,t}^q \\ x_{a,t}(\omega) &= z(\omega)^{\sigma-1} \nu_t^\sigma \kappa_t^{1-\sigma} \phi_{a,t}^x \\ V_{a,t}(\omega) &= z(\omega)^{\sigma-1} \nu_t^\sigma \kappa_t^{1-\sigma} \phi_{a,t}^V. \end{aligned}$$

□

A.2 Proof of Proposition 2

Proposition 2 (Steady State). *In steady state, the share of d 's total expenditure allocated to o is*

$$\pi_{od}^{SS} = \frac{(w_o \tau_{od})^{-\theta} \bar{\phi}_{od}^{ext} \bar{\phi}_{od}^{int}}{\sum_m (w_m \tau_{md})^{-\theta} \bar{\phi}_{md}^{ext} \bar{\phi}_{md}^{int}}, \quad (31)$$

where $\bar{\phi}_{od}^{ext} \equiv \left(\phi_{0,od}^U\right)^{\theta/(\sigma-1)-1}$ and $\bar{\phi}_{od}^{int} \equiv \sum_{a=0}^{\infty} s_{a,od} \rho_{a,od} \left(\phi_{a,od}^q\right)^{(\sigma-1)/\sigma}$.

Proof. Define $\kappa_{od} \equiv w_o \tau_{od}$ as the non-stochastic component of bilateral trade costs. Final expenditure in d on good ω from origin o is given by $x_{od}(\omega) = \nu_d \chi_{od}(\omega) q_{od}^{(\sigma-1)/\sigma} 1\{o = \arg \max_m U_{0,md}\}$.

Total expenditure, $X_{od} = \int_0^1 x_{od}(\omega) d\omega$, is therefore given by,

$$\begin{aligned} X_{od} &= \nu_d^\sigma \kappa_{od}^{1-\sigma} \int_0^1 \chi_{a(\omega),od}(\omega) z_o(\omega)^{\sigma-1} (\phi_{a(\omega),od}^q)^{\frac{\sigma-1}{\sigma}} 1\{c_{od}(\omega)^{1-\sigma} \phi_{0,od}^U = \max_m c_{md}(\omega)^{1-\sigma} \phi_{0,md}^U\} d\omega \\ &= \nu_d^\sigma \kappa_{od}^{1-\sigma} \left(\sum_a s_{a,od} \rho_{a,od} (\phi_{a,od}^q)^{\frac{\sigma-1}{\sigma}} \right) \int_0^1 z_o(\omega)^{\sigma-1} 1\{c_{od}(\omega) (\phi_{0,od}^U)^{\frac{1}{1-\sigma}} = \min_m c_{md}(\omega) (\phi_{0,md}^U)^{\frac{1}{1-\sigma}}\} d\omega \end{aligned}$$

where the second line uses the observation that relationship age is independent of $z_o(\omega)$. Define the shorthand $A_o \equiv \{o \text{ is cheapest}\}$ and write the integral term as $\mathbb{E}[z_o^{\sigma-1} 1\{A_o\}]$. By the Law of Iterated Expectations,

$$\mathbb{E}[z_o^{\sigma-1} 1\{A_o\}] = \mathbb{E}[\mathbb{E}[z_o^{\sigma-1} 1\{A_o\} | z_o = z]] = \mathbb{E}[z_o^{\sigma-1} \mathbb{E}[1\{A_o\} | z_o = z]] = \mathbb{E}[z_o^{\sigma-1} \Pr(A_o | z_o = z)].$$

As draws are independent, some algebra gives

$$\int z^{\sigma-1} f_o(z) \prod_{m \neq o} \Pr\left(z_m \leq \frac{\kappa_{md}}{\kappa_{od}} \left(\frac{\phi_{0,md}^U}{\phi_{0,od}^U}\right)^{\frac{1}{1-\sigma}} z\right) dz = \int z^{\sigma-1} \theta z^{-1-\theta} e^{-(\kappa_{od}(\phi_{0,od}^U)^{\frac{1}{1-\sigma}})^\theta z^{-\theta} \Phi_d} dz,$$

where $\Phi_d \equiv \sum_m (\kappa_{md}(\phi_{0,md}^U)^{\frac{1}{1-\sigma}})^{-\theta}$. Define $t \equiv (\kappa_{od}(\phi_{0,od}^U)^{\frac{1}{1-\sigma}})^\theta \Phi_d z^{-\theta}$, which implies $dz = -(\kappa_{od}(\phi_{0,od}^U)^{\frac{1}{1-\sigma}})^{-\theta} \Phi_d^{-1} \theta^{-1} z^{\theta+1} dt$. Substituting in and using the definition of the Gamma function gives

$$\int_0^\infty t^{-\frac{\sigma-1}{\theta}} \left(\kappa_{od}(\phi_{0,od}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \Phi_d^{\frac{\sigma-1}{\theta}-1} e^{-t} dt = \left(\kappa_{od}(\phi_{0,od}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \Phi_d^{\frac{\sigma-1}{\theta}-1} \Gamma\left(1 - \frac{\sigma-1}{\theta}\right).$$

Substituting this back into the expression for X_{od} gives

$$\begin{aligned} X_{od} &= \nu_d^\sigma \kappa_{od}^{1-\sigma} \left(\sum_a s_{a,od} \rho_{a,od} (\phi_{a,od}^q)^{\frac{\sigma-1}{\sigma}} \right) \left(\kappa_{od}(\phi_{0,od}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \Phi_d^{\frac{\sigma-1}{\theta}-1} \Gamma\left(1 - \frac{\sigma-1}{\theta}\right) \\ &= \kappa_{od}^{-\theta} (\phi_{0,od}^U)^{\frac{\theta}{\sigma-1}-1} \left(\sum_a s_{a,od} \rho_{a,od} (\phi_{a,od}^q)^{\frac{\sigma-1}{\sigma}} \right) \nu_d^\sigma \Phi_d^{\frac{\sigma-1}{\theta}-1} \Gamma\left(1 - \frac{\sigma-1}{\theta}\right). \end{aligned} \quad (32)$$

Defining $\bar{\phi}_{od}^{\text{int}} \equiv \left(\sum_a s_{a,od} \rho_{a,od} (\phi_{a,od}^q)^{\frac{\sigma-1}{\sigma}} \right)$, the trade share is then given by

$$\pi_{od}^{SS} \equiv \frac{X_{od}}{\sum_m X_{md}} = \frac{\kappa_{od}^{-\theta} \bar{\phi}_{od}^{\text{int}} (\phi_{0,od}^U)^{\frac{\theta}{\sigma-1}-1}}{\sum_m \kappa_{md}^{-\theta} \bar{\phi}_{md}^{\text{int}} (\phi_{0,md}^U)^{\frac{\theta}{\sigma-1}-1}}. \quad (33)$$

Finally, the choice probability is simply the integral but without the $z^{\sigma-1}$ term. Note that the z

term cancels out after substituting out dz , so we have

$$\pi_{od}^{src} \equiv \Pr(d \text{ buys from } o) = \int_0^\infty \left(\kappa_{od} (\phi_{0,od}^U)^{\frac{1}{1-\sigma}} \right)^{-\theta} \Phi_d^{-1} e^{-t} dt = \frac{\kappa_{od}^{-\theta} \left(\phi_{0,od}^U \right)^{\frac{\theta}{\sigma-1}}}{\sum_m \kappa_{md}^{-\theta} \left(\phi_{0,md}^U \right)^{\frac{\theta}{\sigma-1}}}. \quad (34)$$

While π_{od}^{src} is not directly observable in data, it is a simple transformation of the (observable) expenditure share. Define $\tilde{\pi}_{od} \equiv \pi_{od} \phi_{0,od}^U / \bar{\phi}_{od}^{\text{int}}$. Then, we have

$$\pi_{od}^{src} = \frac{\tilde{\pi}_{od}}{\sum_m \tilde{\pi}_{md}}. \quad (35)$$

□

A.3 Proof of Proposition 3

Proposition 3 (Transition Path in Relative Changes). *Denote relative changes in variable x_t as $\hat{x}_t \equiv x_t/x_{t-1}$. Consider an unanticipated shock announced in period $\bar{t}$ in relative changes to trade costs in an economy initially in steady state. For a given path of general equilibrium variables expressed in relative changes, $(\hat{w}_{o,t}, \hat{v}_{dj,t}, \hat{P}_{dj,t})$, values for steady state trade shares, $\{\pi_{odj}\}$, and values for the parameters $(\beta, \delta, \sigma, \theta, \{\gamma\}, \mu, \lambda, \{\xi\})$, the path of aggregate bilateral trade flows, $\{\pi_{odj, \bar{t}+s}\}$ for $s \geq 0$, can be calculated without reference to the levels of bilateral trade shifters $\{\tau_{odj}\}$ or of the general equilibrium variables.*

Proof. In what follows, if a quantity that can vary over time does not have a time index, it should be understood as the steady state quantity. Within a period, timing is as follows: (a) any currently unmatched firms form a new match, (b) orders are placed, (c) orders are realized, (d) the technology obsolescence shock happens. The steady state shares, $s_{a,od}$, therefore refer to timing after (a) has occurred. In particular, these follow the recursion $s_{a,od} = \delta \rho_{a-1,od} s_{a,od}$ for $a \geq 1$, with $s_{0,od} = (1 + \sum_{a=1}^\infty \delta^a \prod_{e=0}^{a-1} \rho_{e,od})^{-1}$.

In what follows, some regions for some origins may be empty (e.g., if κ_{od} goes up for some o , there will not be firms switching from o' to o). To avoid notational clutter, we do not include additional min and max functions to account for this on integration limits, so it should be understood that any integrals of the form $\int_L^R$ are zero if $L \geq R$.

Step 1: Derive the distribution of $r_{id,\bar{t}}(\omega)$ Define $b_{d,t}(\omega) \equiv \arg \max_o U_{0,od,t}(\omega)$ as best origin for ω in d in the first period of the shock, $\bar{t}$, if the importer were to choose the origin from scratch (i.e., absent an existing relationship). We will suppress the d, t subscript going forward as it is clear which destination and period are relevant. Then, define

$$r_{id,t}(\omega) \equiv \frac{U_{0,id,t}}{U_{0,bd,t}} \in \left[\frac{\hat{U}_{0,id,t}}{\hat{U}_{0,bd,t}}, 1 \right], \quad (36)$$

as the ratio of the incumbent origin's NPV to the best origin's NPV. Region 1 denotes ω for which $b(\omega) = i$, i.e., the incumbent origin is still the best origin. It will be useful to derive $\Pr(r_{id,\bar{t}}(\omega) \geq r \text{ and } r_{id,\bar{t}-1}(\omega) = 1 | z_i = z)$, i.e., the probability that the post-shock ratio is above some r and i was the incumbent. The probability of the first event is

$$\Pr(r_{id,\bar{t}}(\omega) \geq r | z_i = z) = \Pr(U_{0,id,\bar{t}} \geq r U_{0,bd,\bar{t}}) = \Pr \left(z_m \leq r^{\frac{1}{1-\sigma}} \frac{\kappa_{md,\bar{t}}}{\kappa_{id,\bar{t}}} \left(\frac{\phi_{0,md,\bar{t}}^U}{\phi_{0,id,\bar{t}}^U} \right)^{\frac{1}{1-\sigma}} z \right) \text{ for all } m \neq i.$$

Since the second event is the same but with $r = 1$ and time subscripts $\bar{t} - 1$, the joint probability is

$$\begin{aligned} & f_i(z) \Pr(r_{id,\bar{t}}(\omega) \geq r \text{ and } r_{id,\bar{t}-1}(\omega) = 1 | z_i = z) \\ &= f_i(z) \prod_{m \neq i} \Pr \left(z_m \leq \min \left\{ r^{\frac{1}{1-\sigma}} \frac{\kappa_{md,\bar{t}}}{\kappa_{id,\bar{t}}} \left(\frac{\phi_{0,md,\bar{t}}^U}{\phi_{0,id,\bar{t}}^U} \right)^{\frac{1}{1-\sigma}} z, \frac{\kappa_{md,\bar{t}-1}}{\kappa_{id,\bar{t}-1}} \left(\frac{\phi_{0,md,\bar{t}-1}^U}{\phi_{0,id,\bar{t}-1}^U} \right)^{\frac{1}{1-\sigma}} z \right\} \right) \\ &= \theta z^{-\theta-1} \exp(-z^{-\theta}) \prod_{m \neq i} \Pr \left(z_m \leq z \frac{\kappa_{md,\bar{t}-1}}{\kappa_{id,\bar{t}-1}} \left(\frac{\phi_{0,md,\bar{t}-1}^U}{\phi_{0,id,\bar{t}-1}^U} \right)^{\frac{1}{1-\sigma}} \min \left\{ \left(r \frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{1}{1-\sigma}}, 1 \right\} \right) \\ &= \theta z^{-\theta-1} \exp \left(-z^{-\theta} \kappa_{id,\bar{t}-1}^{\theta} (\phi_{id,\bar{t}-1}^U)^{\frac{\theta}{1-\sigma}} \Phi_{id}(r) \right), \end{aligned}$$

where

$$\Phi_{id}(r) \equiv \sum_m \kappa_{md,\bar{t}-1}^{-\theta} (\phi_{0,md,\bar{t}-1}^U)^{-\frac{\theta}{1-\sigma}} \max \left\{ 1, \left(r \frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \right\}, \quad (37)$$

is the new aggregator. While $\Phi_{id}(r)$ depends on levels of $\kappa_{od,\bar{t}-1}$, note that these are removed when taking hats relative to the steady state aggregator, Φ_d :

$$\frac{\Phi_{id}(r)}{\Phi_d} = \sum_m \pi_{md}^{src} \max \left\{ 1, \left(r \frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \right\}. \quad (38)$$

It will also be useful to calculate $\Phi'_{id}(r)$, which is given by

$$\Phi'_{id}(r) = \frac{\theta}{\sigma-1} \frac{1}{r} \sum_m \kappa_{md,\bar{t}-1}^{-\theta} (\phi_{0,md,\bar{t}-1}^U)^{-\frac{\theta}{1-\sigma}} \left(r \frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} 1 \left\{ r \geq \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}} \right\}. \quad (39)$$

Again, it is useful to normalize it relative to the steady state aggregator:

$$\frac{\Phi'_{id}(r)}{\Phi_d} = \frac{\theta}{\sigma-1} \frac{1}{r} \sum_m \pi_{md}^{src} \left(r \frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} 1 \left\{ r \geq \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}} \right\}. \quad (40)$$

We will repeatedly need to integrate this $z^{\sigma-1}$ over this density, which gives

$$\begin{aligned}
& \mathbb{E}[z_i(\omega)^{\sigma-1} 1\{r_{id,\bar{t}}(\omega) \geq r \text{ and } r_{id,\bar{t}-1}(\omega) = 1\}] \\
&= \int z^{\sigma-1} \theta z^{-\theta-1} \exp\left(-z^{-\theta} \kappa_{id,\bar{t}-1}^{\theta} (\phi_{0,id,\bar{t}-1}^U)^{\frac{\theta}{1-\sigma}} \Phi_{id}(r)\right) dz \\
&= \Gamma\left(1 - \frac{\sigma-1}{\theta}\right) \left(\kappa_{id,\bar{t}-1} (\phi_{0,id,\bar{t}-1}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \Phi_{id}(r)^{\frac{\sigma-1}{\theta}-1}. \quad (41)
\end{aligned}$$

Step 2: Derive expenditure in Region 1 and Region 2 Region 1 is the set of ω for which the best origin is the incumbent origin, i.e., $r_{id,\bar{t}}(\omega) = 1$. Region 2 is the set of ω for which the best origin is no longer the incumbent, but the importer remains with their incumbent supplier because of accumulated relationship capital, i.e., $r_{id,\bar{t}}(\omega) \in [\max\{R_{a,id,\bar{t}}, \hat{U}_{0,id,\bar{t}}/\hat{U}_{0,bd,\bar{t}}\}, 1)$ where $R_{a,id,\bar{t}}$ is the value of $r_{id,\bar{t}}$ at which an importer in a relationship of age a is indifferent between starting fresh in a new origin and remaining with their incumbent supplier. In both cases, the previous contract must be renegotiated, which will in general be a function of both $r_{id,\bar{t}}(\omega)$ (as this governs how much better the new-best origin is for good ω) and $b(\omega)$ (as the contract will depend on the full time-path of the outside option, which is origin specific), so we write $\phi_{a,od,\bar{t}}^q(r, m)$ to make this explicit.

We exclude from Region 1 any relationships that broke up at the end of $\bar{t}-1$ but choose to remain in the same origin, as these contracts do not depend on the pre-shock period in any way. We will account for these among the new “ordinary” relationship stock. Following the same steps as in Proposition 2, aggregate expenditure in $\bar{t}$ for goods in Region 1 is

$$\begin{aligned}
X_{id,\bar{t}}^{R1} &= \nu_{d,\bar{t}}^{\sigma} \kappa_{id,\bar{t}}^{1-\sigma} \left(\sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} (\phi_{a,id,\bar{t}}^q(1))^{\frac{\sigma-1}{\sigma}} \right) \int z^{\sigma-1} f_i(z) \Pr(r_{id,\bar{t}}(\omega) = 1 \text{ and } r_{id,\bar{t}-1}(\omega) = 1 | z_i = z) dz \\
&= \nu_{d,\bar{t}}^{\sigma} \kappa_{id,\bar{t}}^{1-\sigma} \left(\sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} (\phi_{a,id,\bar{t}}^q(1))^{\frac{\sigma-1}{\sigma}} \right) \Gamma\left(1 - \frac{\sigma-1}{\theta}\right) \left(\kappa_{id,\bar{t}-1} (\phi_{0,id,\bar{t}-1}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \Phi_{id}(1)^{\frac{\sigma-1}{\theta}-1}.
\end{aligned}$$

where the second line uses (41). Dividing by (32) gives

$$\begin{aligned}
\frac{X_{id,\bar{t}}^{R1}}{X_{id,\bar{t}-1}} &= \hat{\nu}_{d,\bar{t}}^{\sigma} \hat{\kappa}_{id,\bar{t}}^{1-\sigma} \left(\frac{\sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} (\phi_{a,id,\bar{t}}^q(1))^{\frac{\sigma-1}{\sigma}}}{\bar{\phi}_{id}^{\text{int}}} \right) \left(\frac{\Phi_{id}(1)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \\
&= \hat{\nu}_{d,\bar{t}}^{\sigma} \hat{\kappa}_{id,\bar{t}}^{1-\sigma} \left(\frac{\sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} (\phi_{a,id,\bar{t}}^q(1))^{\frac{\sigma-1}{\sigma}}}{\bar{\phi}_{id}^{\text{int}}} \right) \left(\sum_m \pi_{md}^{\text{src}} \max \left\{ 1, \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \right\} \right)^{\frac{\sigma-1}{\theta}-1}, \quad (42)
\end{aligned}$$

which does not depend on levels.

In Region 2, the renegotiated contract depends on the level of $r_{id,\bar{t}}(\omega)$ and on the identity of the the new best origin as different origins have different wage dynamics. We therefore need to integrate over the joint density of $(r_{id,\bar{t}}(\omega), b(\omega))$. Differentiating (41) to get the relevant expenditure-

weighted density of productivity at a given r gives

$$\begin{aligned} & -\frac{d}{dr} \mathbb{E}[z_i(\omega)^{\sigma-1} 1\{r_{id,\bar{t}}(\omega) \geq r \text{ and } r_{id,\bar{t}-1}(\omega) = 1\}] \\ & = \Gamma\left(1 - \frac{\sigma-1}{\theta}\right) \left(\kappa_{id,\bar{t}-1}(\phi_{0,id,\bar{t}-1}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \left(1 - \frac{\sigma-1}{\theta}\right) \Phi'_{id}(r) \Phi_{id}(r)^{\frac{\sigma-1}{\theta}-2}. \end{aligned}$$

This is the $z_i^{\sigma-1}$ -weighted density for the event that any origin $m \neq i$ is the new best. To obtain the density for the portion that a particular m is the new best, we can simply take the portion of (39) for a given m as the marginal goods affected by moving r must be those for which the origin in (39) is the new best (since ties occur with probability zero). The relevant density is therefore

$$\begin{aligned} g_{id,\bar{t}}(r, m) & \equiv -\frac{d}{dr} \mathbb{E}[z_i(\omega)^{\sigma-1} 1\{r_{id,\bar{t}}(\omega) \geq r \text{ and } r_{id,\bar{t}-1}(\omega) = 1 \text{ and } m \text{ is the new best origin}\}] \\ & = \Gamma\left(1 - \frac{\sigma-1}{\theta}\right) \left(\kappa_{id,\bar{t}-1}(\phi_{0,id,\bar{t}-1}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \left(1 - \frac{\sigma-1}{\theta}\right) \Phi_{id}(r)^{\frac{\sigma-1}{\theta}-2} \\ & \quad \times \frac{\theta}{\sigma-1} \frac{1}{r} \kappa_{md,\bar{t}-1}^{-\theta} (\phi_{0,md,\bar{t}-1}^U)^{-\frac{\theta}{1-\sigma}} \left(r \frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}}\right)^{\frac{\theta}{\sigma-1}} 1\left\{r \geq \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}\right\}. \end{aligned}$$

Then, expenditure in Region 2 is given by

$$\begin{aligned} X_{id,\bar{t}}^{R2} & = \nu_{d,\bar{t}}^\sigma \kappa_{id,\bar{t}}^{1-\sigma} \sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} \sum_{m \neq i} \int_{\max\left\{R_{a,imd,\bar{t}}, \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}\right\}}^1 (\phi_{a,id,\bar{t}}^q(r, m))^{\frac{\sigma-1}{\sigma}} g_{id,\bar{t}}(r, m) dr \\ & = \nu_{d,\bar{t}}^\sigma \kappa_{id,\bar{t}}^{1-\sigma} \Gamma\left(1 - \frac{\sigma-1}{\theta}\right) \left(\kappa_{id,\bar{t}-1}(\phi_{0,id,\bar{t}-1}^U)^{\frac{1}{1-\sigma}}\right)^{\sigma-1-\theta} \left(\frac{\theta}{\sigma-1} - 1\right) \sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} \\ & \quad \times \sum_{m \neq i} \kappa_{md,\bar{t}-1}^{-\theta} (\phi_{0,md,\bar{t}-1}^U)^{-\frac{\theta}{1-\sigma}} \int_{\max\left\{R_{a,imd,\bar{t}}, \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}\right\}}^1 (\phi_{a,id,\bar{t}}^q(r, m))^{\frac{\sigma-1}{\sigma}} \frac{1}{r} \left(r \frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}}\right)^{\frac{\theta}{\sigma-1}} \Phi_{id}(r)^{\frac{\sigma-1}{\theta}-2} dr. \end{aligned}$$

This simplifies considerably when dividing by (32) to obtain the change. In particular,

$$\begin{aligned} \frac{X_{id,\bar{t}}^{R2}}{X_{id,\bar{t}-1}} & = \hat{\nu}_{d,\bar{t}}^\sigma \hat{\kappa}_{id,\bar{t}}^{1-\sigma} \frac{1}{\bar{\phi}_{id}^{\text{int}}} \left(\frac{\theta}{\sigma-1} - 1\right) \sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} \sum_{m \neq i} \frac{\kappa_{md,\bar{t}-1}^{-\theta} (\phi_{0,md,\bar{t}-1}^U)^{-\frac{\theta}{1-\sigma}}}{\Phi_d} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}}\right)^{\frac{\theta}{\sigma-1}} \\ & \quad \times \int_{\max\left\{R_{a,imd,\bar{t}}, \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}\right\}}^1 (\phi_{a,id,\bar{t}}^q(r, m))^{\frac{\sigma-1}{\sigma}} r^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(r)}{\Phi_d}\right)^{\frac{\sigma-1}{\theta}-2} dr \\ & = \hat{\nu}_{d,\bar{t}}^\sigma \hat{\kappa}_{id,\bar{t}}^{1-\sigma} \frac{1}{\bar{\phi}_{id}^{\text{int}}} \left(\frac{\theta}{\sigma-1} - 1\right) \sum_{a=1}^{\infty} s_{a,id} \rho_{a,id} \sum_{m \neq i} \pi_{md}^{\text{src}} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}}\right)^{\frac{\theta}{\sigma-1}} \\ & \quad \times \int_{\max\left\{R_{a,imd,\bar{t}}, \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}\right\}}^1 (\phi_{a,id,\bar{t}}^q(r, m))^{\frac{\sigma-1}{\sigma}} r^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(r)}{\Phi_d}\right)^{\frac{\sigma-1}{\theta}-2} dr. \end{aligned} \tag{43}$$

Recalling that $\Phi_{id}(r)/\Phi_d$ is given by (38) and does not depend on levels, we can see that this overall expression does not depend on levels.

Step 3: Derive expenditure in Region 3 and exogenous resets In Region 3, the new best origin is sufficiently better than the incumbent origin that the importer is willing to forgo their current supplier and switch immediately. This is characterized by the region $r \in [\hat{U}_{0,id,\bar{t}}/\hat{U}_{0,md,\bar{t}}, R_{imd,\bar{t}}]$. Unlike Region 2, the contract itself does not vary with r , as once the switch happened the outside option is to stay in the new origin. However, the mass of Region 3 firms does depend on the distribution of r . We can re-use the z_i -expenditure-weighted density of r from Region 2 by substituting out $z_m^{\sigma-1} r_{id,t}(\omega)$ in favor of $z_i^{\sigma-1}$ by re-arranging the definition of $r_{id,\bar{t}}(\omega)$ (36) to obtain $\kappa_{md,\bar{t}}^{1-\sigma} z_m^{\sigma-1} = \kappa_{id,\bar{t}}^{1-\sigma} z_i^{\sigma-1} \phi_{0,id,\bar{t}}^U / (r \phi_{0,md,\bar{t}}^U)$. Expenditure is therefore given by

$$\begin{aligned} X_{i \rightarrow m,d,\bar{t}}^{R3} &= \nu_{d,\bar{t}}^\sigma \kappa_{id,\bar{t}}^{1-\sigma} \left(\frac{\phi_{0,id,\bar{t}}^U}{\phi_{0,md,\bar{t}}^U} \right) \rho_{0,md}(\phi_{0,md,\bar{t}}^q)^{\frac{\sigma-1}{\sigma}} \sum_{a=0}^{\infty} \tilde{s}_{a,id} \int_{\frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}}^{R_{a,imd,\bar{t}}} \frac{1}{r} g_{id,\bar{t}}(r, m) dr \\ &= \nu_{d,\bar{t}}^\sigma \kappa_{id,\bar{t}}^{1-\sigma} \left(\frac{\phi_{0,id,\bar{t}}^U}{\phi_{0,md,\bar{t}}^U} \right) \rho_{0,md}(\phi_{0,md,\bar{t}}^q)^{\frac{\sigma-1}{\sigma}} \Gamma\left(1 - \frac{\sigma-1}{\theta}\right) \left(\kappa_{id,\bar{t}-1}(\phi_{0,id,\bar{t}-1}^U)^{\frac{1}{1-\sigma}} \right)^{\sigma-1-\theta} \left(\frac{\theta}{\sigma-1} - 1 \right) \\ &\quad \times \kappa_{md,\bar{t}-1}^{-\theta} (\phi_{0,md,\bar{t}-1}^U)^{-\frac{\theta}{1-\sigma}} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \sum_{a=0}^{\infty} \tilde{s}_{a,id} \int_{\frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}}^{R_{a,imd,\bar{t}}} r^{\frac{\theta}{\sigma-1}-2} \Phi_{id}(r)^{\frac{\sigma-1}{\theta}-2} dr, \end{aligned}$$

where $\tilde{s}_{a,id} = s_{0,id} - (1 - \delta)$ for $a = 0$ and $\tilde{s}_{a,id} = s_{a,id}$ for $a \geq 1$. As we substituted in $z_i^{\sigma-1}$, to obtain all terms in hats we need to divide by $X_{id,\bar{t}-1}$. However, we can obtain the object of interest, $X_{i \rightarrow m,d,\bar{t}}^{R3}/X_{md,\bar{t}-1}$, as follows. Since total Region 3 inflows into an origin m are given by $X_{md,\bar{t}}^{R3} = \sum_{i \neq m} X_{i \rightarrow m,d,\bar{t}}^{R3}$, we can write

$$\frac{X_{md,\bar{t}}^{R3}}{X_{md,\bar{t}-1}} = \sum_{i \neq m} \frac{X_{id,\bar{t}-1}}{X_{md,\bar{t}-1}} \frac{X_{i \rightarrow m,d,\bar{t}}^{R3}}{X_{id,\bar{t}-1}} = \sum_{i \neq m} \frac{\pi_{id,\bar{t}-1}}{\pi_{md,\bar{t}-1}} \frac{X_{i \rightarrow m,d,\bar{t}}^{R3}}{X_{id,\bar{t}-1}}. \quad (44)$$

Then, following the same steps as in Region 2, we have

$$\begin{aligned} \frac{X_{md,\bar{t}}^{R3}}{X_{md,\bar{t}-1}} &= \sum_{i \neq m} \frac{\pi_{id,\bar{t}-1}}{\pi_{md,\bar{t}-1}} \hat{\nu}_{d,\bar{t}}^\sigma \hat{\kappa}_{id,\bar{t}}^{1-\sigma} \left(\frac{\phi_{0,id,\bar{t}}^U}{\phi_{0,md,\bar{t}}^U} \right) \rho_{0,md} \frac{1}{\phi_{id}^{\text{int}}} (\phi_{0,md,\bar{t}}^q)^{\frac{\sigma-1}{\sigma}} \left(\frac{\theta}{\sigma-1} - 1 \right) \\ &\quad \times \pi_{md}^{\text{src}} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \sum_{a=0}^{\infty} \tilde{s}_{a,id} \int_{\frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}}^{R_{a,imd,\bar{t}}} r^{\frac{\theta}{\sigma-1}-2} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-2} dr \\ &= \hat{\nu}_{d,\bar{t}}^\sigma \hat{\kappa}_{md,\bar{t}}^{1-\sigma} \rho_{0,md} \frac{1}{\phi_{md}^{\text{int}}} (\phi_{0,md,\bar{t}}^q)^{\frac{\sigma-1}{\sigma}} \left(\frac{\theta}{\sigma-1} - 1 \right) \\ &\quad \times \sum_{i \neq m} \pi_{id}^{\text{src}} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}-1} \sum_{a=0}^{\infty} \tilde{s}_{a,id} \int_{\frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}}^{R_{a,imd,\bar{t}}} r^{\frac{\theta}{\sigma-1}-2} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-2} dr. \quad (45) \end{aligned}$$

Moreover, we can obtain the integral in closed-form. Recall from (38) that $\Phi_{id}(r)/\Phi_d$ is given by $\sum_l \pi_{ld}^{src} \max\{1, (r\hat{U}_{0,ld,\bar{t}}/\hat{U}_{0,id,\bar{t}})^{\theta/(\sigma-1)}\}$. Thus, for a given origin l , define

$$r_{lid} \equiv \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,ld,\bar{t}}}, \quad (46)$$

as the lowest value of r for which the max expression takes the r branch (i.e., the kink point). Then, for an integration region $[L, U]$, let $\mathcal{K}_{id}(L, U)$ denote the ordered set consisting of the integration endpoints and all kink points, r_{lid} , lying strictly between them. Formally, $\mathcal{K}_{id}(L, U) \equiv \text{sort}(\{L, U\} \cup \{r_{lid} : L < r_{lid} < U\})$, with ordered elements

$$L = k_{id,0}(L, U) < k_{id,1}(L, U) < \dots < k_{id,K_{id}(L,U)} = U. \quad (47)$$

This creates $K_{id}(L, U)$ subintervals:

$$[k_{id,0}, k_{id,1}], [k_{id,1}, k_{id,2}], \dots, [k_{id,K_{id}(L,U)-1}, k_{id,K_{id}(L,U)}].$$

By construction, the max term never switches branch within any subinterval. For Region 3, we have

$$L = L_{imd}^{R3} \equiv \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}, \quad U = U_{a,imd}^{R3} \equiv R_{a,imd,\bar{t}}. \quad (48)$$

Then, write the elements themselves as $k_{a,imd,h}^{R3} \equiv k_{id,h}(L_{imd}^{R3}, U_{a,imd}^{R3})$ and $K_{a,imd}^{R3} \equiv K_{id}(L_{imd}^{R3}, U_{a,imd}^{R3})$.

For concreteness, consider a given subinterval, $[k_{a,imd,h}^{R3}, k_{a,imd,h+1}^{R3}]$. Some algebra shows that for $r \in [k_{a,imd,h}^{R3}, k_{a,imd,h+1}^{R3}]$,

$$\frac{d}{dr} \left[r^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \right] = \left(\frac{\theta}{\sigma-1} - 1 \right) r^{\frac{\theta}{\sigma-1}-2} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-2} \left[\sum_{l:r_{lid} \geq k_{a,imd,h+1}^{R3}} \pi_{ld}^{src} \right]. \quad (49)$$

Substituting (49) into the integral term in (45) gives

$$\begin{aligned} & \left(\frac{\theta}{\sigma-1} - 1 \right) \int_{k_{a,imd,h}^{R3}}^{k_{a,imd,h+1}^{R3}} r^{\frac{\theta}{\sigma-1}-2} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-2} dr \\ &= \left[\sum_{l:r_{lid} \geq k_{a,imd,h+1}^{R3}} \pi_{ld}^{src} \right]^{-1} \int_{k_{a,imd,h}^{R3}}^{k_{a,imd,h+1}^{R3}} \frac{d}{dr} \left[r^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \right] dr \\ &= \left[\sum_{l:r_{lid} \geq k_{a,imd,h+1}^{R3}} \pi_{ld}^{src} \right]^{-1} \left((k_{a,imd,h+1}^{R3})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,imd,h+1}^{R3})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} - (k_{a,imd,h}^{R3})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,imd,h}^{R3})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \right) \end{aligned}$$

Overall, then, substituting this back into (45) gives

$$\begin{aligned} \frac{X_{md,\bar{t}}^{R3}}{X_{md,\bar{t}-1}^{R3}} &= \hat{\nu}_{d,\bar{t}}^\sigma \hat{\kappa}_{md,\bar{t}}^{1-\sigma} \rho_{0,md} \frac{1}{\bar{\phi}_{md}^{\text{int}}} (\phi_{0,md,\bar{t}}^q)^{\frac{\sigma-1}{\sigma}} \sum_{i \neq m} \pi_{id}^{\text{src}} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}-1} \sum_{a=0}^{\infty} \tilde{s}_{a,id} \sum_{h=0}^{K_{a,imd}^{R3}-1} \left[\sum_{l: r_{lid} \geq k_{a,imd,h+1}^{R3}} \pi_{ld}^{\text{src}} \right]^{-1} \\ &\times \left((k_{a,imd,h+1}^{R3})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,imd,h+1}^{R3})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} - (k_{a,imd,h}^{R3})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,imd,h}^{R3})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \right). \end{aligned} \quad (50)$$

Step 4: Iterating forward and accounting We have now derived the expenditure within each region for the period of the shock, $\bar{t}$. Region 1 and Region 2 are legacy relationships in the sense that the contracts are renegotiated versions of pre-shock contracts, and both groups will decay towards zero as relationships break up over time. All other goods start with new contracts in the shock period, which have no dependence on pre-shock information. Thus, we can decompose total expenditure in a given period $t \geq \bar{t}$ as $X_{od,t} = X_{od,t}^{R1} + X_{od,t}^{R2} + X_{od,t}^N$.

We first extend the Region 1 and Region 2 formulas to be for a generic period $\bar{t} + j$ with $j \geq 0$. For Region 1, it is straightforward to extend (42) as follows:

$$\frac{X_{id,\bar{t}+j}^{R1}}{X_{id,\bar{t}-1}^{R1}} = \hat{\nu}_{d,\bar{t}+j}^\sigma \hat{\kappa}_{id,\bar{t}+j}^{1-\sigma} \left(\frac{\sum_{a=1+j}^{\infty} s_{a,id} \rho_{a,id} (\phi_{a,id,\bar{t}}^q)^{\frac{\sigma-1}{\sigma}}}{\bar{\phi}_{id}^{\text{int}}} \right) \left(\sum_m \pi_{md}^{\text{src}} \max \left\{ 1, \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \right\} \right)^{\frac{\sigma-1}{\theta}-1},$$

where the utility ratio term inside the max does not increase with j because the terms correspond to the probability of being in Region 1, which is determined at $\bar{t}$. For Region 2, we can extend (??) as follows:

$$\begin{aligned} \frac{X_{id,\bar{t}+j}^{R2}}{X_{id,\bar{t}-1}^{R2}} &= \hat{\nu}_{d,\bar{t}+j}^\sigma \hat{\kappa}_{id,\bar{t}+j}^{1-\sigma} \frac{1}{\bar{\phi}_{id}^{\text{int}}} \left(\frac{\theta}{\sigma-1} - 1 \right) \sum_{a=1+j}^{\infty} s_{a,id} \rho_{a,id} \sum_{m \neq i} \pi_{md}^{\text{src}} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \\ &\times \int_{\max \left\{ R_{a-j,imd,\bar{t}}, \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,md,\bar{t}}} \right\}}^1 (\phi_{a,id,\bar{t}}^q(r, m))^{\frac{\sigma-1}{\sigma}} r^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-2} dr, \end{aligned}$$

where recall that $\phi_{a,id,\bar{t}}^q(r, m)$ refers to the quantity wedge in relationship age a in a contract signed (re-negotiated) in period $\bar{t}$ where the incumbent origin was i and when the ratio $r_{imd,\bar{t}}(\omega) = r$ in period $\bar{t}$ and the best origin at the date of signing was m .

Define the productivity-weighted mass of ω in an ordinary relationship of age a at time t as

$$G_{a,od,t} \equiv \frac{\mathbb{E}[z_o(\omega)^{\sigma-1} 1\{\omega \text{ is in an ordinary relationship in } o \text{ of age } a \text{ at } t\}]}{\mathbb{E}[z_o(\omega)^{\sigma-1} 1\{o \text{ is the incumbent origin for } \omega \text{ at } \bar{t}-1\}]}, \quad (51)$$

where “ordinary” means that the relationship was initiated on or after $\bar{t}$. Then, by construction,

ordinary-pool expenditure relative to steady state expenditure is

$$\frac{X_{od,\bar{t}+j}^G}{X_{od,\bar{t}-1}} = \hat{\nu}_{d,\bar{t}+j}^\sigma \hat{\kappa}_{od,\bar{t}+j}^{1-\sigma} \frac{1}{\bar{\phi}_{od}^{\text{int}}} \sum_{a=0}^j \rho_{a,od} \left(\phi_{a,od,\bar{t}+j-a}^q \right)^{\frac{\sigma-1}{\sigma}} G_{a,od,\bar{t}+j}. \quad (52)$$

We will compute $G_{a,od,t}$ recursively. $G_{a,od,t}$ evolves according to the following rule:

$$G_{a,od,t+1} = \begin{cases} G_{0,od,t+1}^{exo} + G_{0,od,t+1}^{G \rightarrow G} + G_{0,od,t+1}^{R1 \rightarrow G} + G_{0,od,t+1}^{R2(FB) \rightarrow G} + G_{0,od,t+1}^{R2(remain) \rightarrow G} & \text{for } a = 0 \\ \delta \rho_{a-1,od} G_{a-1,od,t} & \text{for } a \geq 1, \end{cases} \quad (53)$$

where $G_{0,od,t+1}^{exo}$ is the mass flowing into o from relationships that broke up exogenously at the end of period t , $G_{0,od,t+1}^{G \rightarrow G}$, $G_{0,od,t+1}^{R1 \rightarrow G}$, $G_{0,od,t+1}^{R2(FB) \rightarrow G}$, $G_{0,od,t+1}^{R2(remain) \rightarrow G}$ are the masses flowing into o from relationships that broke up through failure at the end of period t among ordinary relationships, Region 1 relationships, Region 2 relationships for whom o is the fallback, and Region 2 relationships for whom o is the incumbent, respectively.

To initialize $G_{a,od,t}$, note that for $t = \bar{t}$, $G_{a,od,t} = 0$ for $a \geq 1$, as all pre-existing relationships are in Region 1 or Region 2. We thus only need to initialize $G_{0,od,\bar{t}}$, which is given by Region 3, the exogenous breakups from $\bar{t} - 1$, and the relationships that broke up through failure in $\bar{t} - 1$ but remain in the incumbent origin (those that switched are included in Region 3), i.e., $G_{0,od,\bar{t}} = G_{0,od,\bar{t}}^{R3} + G_{0,od,\bar{t}}^{exo} + G_{0,od,\bar{t}}^{R1-failure}$. From (50), the Region 3 component is

$$G_{0,od,\bar{t}}^{R3} = \sum_{i \neq o} \pi_{id}^{src} \left(\frac{\hat{U}_{0,od,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}-1} \sum_{a=0}^{\infty} \tilde{s}_{a,id} \sum_{h=0}^{K_{a,i,od}^{R3}-1} \left[\sum_{l: \tau_{lid} \geq k_{a,i,od,h+1}^{R3}} \pi_{ld}^{src} \right]^{-1} \\ \times \left((k_{a,i,od,h+1}^{R3})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,i,od,h+1}^{R3})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} - (k_{a,i,od,h}^{R3})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,i,od,h}^{R3})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \right).$$

The exogenous breakup component for a generic j is

$$G_{0,od,\bar{t}+j}^{exo} = (1 - \delta) \left(\hat{\kappa}_{od,\bar{t}+j} (\hat{\phi}_{0,od,\bar{t}+j}^U)^{\frac{1}{1-\sigma}} \right)^{\sigma-1-\theta} \left(\sum_m \pi_{md}^{src} \hat{\kappa}_{md,\bar{t}+j}^{-\theta} (\hat{\phi}_{0,md,\bar{t}+j}^U)^{-\frac{\theta}{1-\sigma}} \right)^{\frac{\sigma-1}{\theta}-1}.$$

The failure-induced breakups component is

$$G_{0,od,\bar{t}}^{R1-failure} = (s_{0,od} - (1 - \delta)) \left(\sum_m \pi_{md}^{src} \max \left\{ 1, \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,od,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \right\} \right)^{\frac{\sigma-1}{\theta}-1}. \quad (54)$$

All the remains is to construct $G_{0,od,t+1}^{G \rightarrow G}$, $G_{0,od,t+1}^{R1 \rightarrow G}$, $G_{0,od,t+1}^{R2 \rightarrow G}$. As ordinary relationships that end

restart in the same origin, we have $G_{0,od,t+1}^{G \rightarrow G} = \delta \sum_{a=0}^{\infty} (1 - \rho_{a,od}) G_{a,od,t}$. For Region 1, we have

$$G_{0,od,\bar{t}+j+1}^{R1 \rightarrow G} = \delta \left(\sum_m \pi_{md}^{src} \max \left\{ 1, \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,od,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \right\} \right)^{\frac{\sigma-1}{\theta}-1} \sum_{a=1+j}^{\infty} s_{a,od} (1 - \rho_{a,od}). \quad (55)$$

For Region 2, relationships with origin i that break up through failure either switch to the new-best origin m or remain with the incumbent origin (but with a new supplier), depending on which is better at the moment of the failure. This is determined by whether $r_{imd,t}(\omega) < 1$ in the period t where the decision is being made. Importantly, this is simply a function of hats relative to the $r_{imd,\bar{t}}(\omega)$, so we can define a cutoff

$$r_{imd,\bar{t}+j+1}^* \equiv \frac{\hat{U}_{0,md,\bar{t}+j+1} / \hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,id,\bar{t}+j+1} / \hat{U}_{0,id,\bar{t}}}, \quad (56)$$

which has the property that the fallback m is chosen if $r < r_{imd,\bar{t}+j+1}^*$ and the incumbent i is chosen if $r \geq r_{imd,\bar{t}+j+1}^*$. For the goods with $r < r_{imd,\bar{t}+j+1}^*$, this is conceptually the same as Region 3 but applied to an age 0 relationship in a future period. So, following the same steps yields

$$\begin{aligned} G_{0,od,\bar{t}+j+1}^{R2(FB) \rightarrow G} &= \delta \sum_{i \neq o} \pi_{id}^{src} \left(\frac{\hat{U}_{0,od,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}-1} \sum_{a=j+1}^{\infty} s_{a,id} (1 - \rho_{a,id}) \\ &\quad \times \left(\frac{\theta}{\sigma-1} - 1 \right) \int_{\max\{R_{a-j,i,od,\bar{t}}, \frac{\hat{U}_{0,id,\bar{t}}}{\hat{U}_{0,od,\bar{t}}}\}}^{\min\{1, r_{iod,\bar{t}+j+1}^*\}} r^{\frac{\theta}{\sigma-1}-2} \left(\frac{\Phi_{id}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-2} dr. \end{aligned}$$

As in the discussion of Region 3, we can obtain this integral in closed form. Define $L_{a,i,od,j}^{R2(FB)} \equiv \max\{R_{a-j,i,od,\bar{t}}, \hat{U}_{0,id,\bar{t}} / \hat{U}_{0,od,\bar{t}}\}$ and $U_{iod,j}^{R2(FB)} \equiv \min\{1, r_{iod,\bar{t}+j+1}^*\}$. Then, define the cutpoints $k_{a,i,od,j,h}^{R2(FB)} \equiv k_{id,h}(L_{a,i,od,j}^{R2(FB)}, U_{iod,j}^{R2(FB)})$ and total number of intervals $K_{a,i,od,j}^{R2(FB)} \equiv K_{id}(L_{a,i,od,j}^{R2(FB)}, U_{iod,j}^{R2(FB)})$. We then have

$$\begin{aligned} G_{0,od,\bar{t}+j+1}^{R2(FB) \rightarrow G} &= \delta \sum_{i \neq o} \pi_{id}^{src} \left(\frac{\hat{U}_{0,od,\bar{t}}}{\hat{U}_{0,id,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}-1} \sum_{a=j+1}^{\infty} s_{a,id} (1 - \rho_{a,id}) \sum_{h=0}^{K_{a,i,od,j}^{R2(FB)}-1} \left[\sum_{l: r_{lid} \geq k_{a,i,od,j,h+1}^{R2(FB)}} \pi_{ld}^{src} \right]^{-1} \\ &\quad \times \left((k_{a,i,od,j,h+1}^{R2(FB)})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,i,od,j,h+1}^{R2(FB)})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} - (k_{a,i,od,j,h}^{R2(FB)})^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{id}(k_{a,i,od,j,h}^{R2(FB)})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \right), \end{aligned}$$

For the goods with $r \geq r_{imd,\bar{t}+j+1}^*$

$$G_{0,od,\bar{t}+j+1}^{R2(remain) \rightarrow G} = \delta \left(\frac{\theta}{\sigma-1} - 1 \right) \sum_{m \neq o} \pi_{md}^{src} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,od,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \sum_{a=1+j}^{\infty} s_{a,od} (1 - \rho_{a,od}) \\ \times \int_{\max \left\{ R_{a-j,od,\bar{t}}, \frac{\hat{U}_{0,od,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}, r_{imd,\bar{t}+j+1}^* \right\}}^1 r^{\frac{\theta}{\sigma-1}-1} \left(\frac{\Phi_{od}(r)}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-2} dr.$$

The integral in this region is similarly available in closed form. Define $L_{a,od,j}^{R2(remain)} \equiv \max \left\{ R_{a-j,od,\bar{t}}, \frac{\hat{U}_{0,od,\bar{t}}}{\hat{U}_{0,md,\bar{t}}}, r_{imd,\bar{t}+j+1}^* \right\}$

Define the cutpoints $k_{a,od,j,h}^{R2(remain)} \equiv k_{od,h}(L_{a,od,j}^{R2(remain)}, 1)$ and number of intervals $K_{a,od,j}^{R2(remain)} \equiv K_{od}(L_{a,od,j}^{R2(remain)}, 1)$. As the exponent on the r term in the integral is different, some algebra gives

$$G_{0,od,\bar{t}+j+1}^{R2(remain) \rightarrow G} = \delta \sum_{m \neq o} \pi_{md}^{src} \left(\frac{\hat{U}_{0,md,\bar{t}}}{\hat{U}_{0,od,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \sum_{a=1+j}^{\infty} s_{a,od} (1 - \rho_{a,od}) \sum_{h=0}^{K_{a,od,j}^{R2(remain)}-1} \left[\sum_{l: r_{lod} \leq k_{a,od,j,h}^{R2(remain)}} \pi_{ld}^{src} \left(\frac{\hat{U}_{0,ld,\bar{t}}}{\hat{U}_{0,od,\bar{t}}} \right)^{\frac{\theta}{\sigma-1}} \right] \\ \times \left(\left(\frac{\Phi_{od}(k_{a,od,j,h}^{R2(remain)})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} - \left(\frac{\Phi_{od}(k_{a,od,j,h+1}^{R2(remain)})}{\Phi_d} \right)^{\frac{\sigma-1}{\theta}-1} \right).$$

Overall, then, the relative change in expenditure is given by

$$\frac{X_{od,\bar{t}+j}}{X_{od,\bar{t}-1}} = \frac{X_{od,\bar{t}+j}^{R1}}{X_{od,\bar{t}-1}} + \frac{X_{od,\bar{t}+j}^{R2}}{X_{od,\bar{t}-1}} + \frac{X_{od,\bar{t}+j}^G}{X_{od,\bar{t}-1}}, \quad (57)$$

with $X_{od,\bar{t}+j}^G/X_{od,\bar{t}-1}$ computed from (52) and $G_{a,od,\bar{t}+j}$ computed from (53).

□

B Scale Equivariance: From Micro Problems to Quantitative Trade

This appendix explains the scale equivariance property and why it is useful beyond the particular relational contracting model studied in the paper. Scale equivariance is a property of a microeconomic decision problem, which may be a contract between two firms, an investment problem faced by one firm, a firm's decision to build a customer base in a destination, or another dynamic problem that a researcher may wish to embed within a quantitative trade model.

It is useful to keep three layers conceptually distinct. First, the microeconomic problem faced by an individual unit. Second, the aggregation of the outcome of (potentially infinitely) many such units that may differ by permanent productivity draws or aggregate economic conditions or microeconomic states. Third, the embedding of those aggregate objects into a trade structure, such as [Eaton and Kortum \(2002\)](#) sourcing or [Melitz \(2003\)](#)-Pareto export participation. Scale equivariance concerns the first layer, but its value is that it can make the second and third layers simpler.

It is also useful to note that scale equivariance is relevant to quantitative trade models in two separate senses. First, it allows researchers to embed a microeconomic problem within a quantitative trade model without having to solve the microeconomic problem for every possible value of heterogeneous productivity (nor to discretize the productivity space, which would sacrifice closed-form aggregation). Second, even if the problem is not at a microeconomic level that requires aggregation, it implies that the current value of economic fundamentals can be factored out of the problem and thus permits a solution using dynamic hat algebra. For example, [Kleinman et al. \(2023\)](#) introduce landlords that invest in capital in each location. While there is no aggregation concern as there is one representative landlord per location, the fact that the capital investment problem satisfies equivariance means that investment decisions depend only on relative changes in economic fundamentals.

B.1 Examples of scale equivariant micro problems

We now provide four examples of scale equivariant micro problems, designed to highlight that this property is not specific to the model of relational contracting presented in this paper (or even contracting more generally).

B.1.1 Example 1: The relational contracting problem in this paper

Consider a given importer-supplier match and suppress all aggregate indices. Assume that the outside option is re-matching at age 0. Let productivity be z , so marginal cost is $c_t = \kappa_t/z$. The importer's revenue is $\nu_t q^{(\sigma-1)/\sigma}$. The only state variable is relationship age a . The contract specifies quantities $q_{a,t}$ and transfers $x_{a,t}$. The buyer's continuation value is $V_{a,t}$.

Suppose productivity increases from z to λz . Consider transforming every quantity, payment, and

continuation value according to

$$q'_{a,t} = \lambda^\sigma q_{a,t}, \quad x'_{a,t} = \lambda^{\sigma-1} x_{a,t}, \quad V'_{a,t} = \lambda^{\sigma-1} V_{a,t}. \quad (58)$$

Observe that both the buyer's and the seller's flow payoffs are simply scaled by $\lambda^{\sigma-1}$:

$$\pi_{a,t}^{b'} = \nu_t (q'_{a,t})^{(\sigma-1)/\sigma} - x'_{a,t} = \nu_t (\lambda^\sigma q_{a,t})^{(\sigma-1)/\sigma} - \lambda^{\sigma-1} x_{a,t} = \lambda^{\sigma-1} \pi_{a,t}^b \quad (59)$$

$$\pi_{a,t}^{s'} = x'_{a,t} - \kappa_t / (\lambda z) q'_{a,t} = \lambda^{\sigma-1} x_{a,t} - \lambda^{\sigma-1} \kappa_t / z q_{a,t} = \lambda^{\sigma-1} \pi_{a,t}^s. \quad (60)$$

The learning process is unchanged as success probabilities and beliefs do not depend on z . The seller's breakeven constraint is unchanged by (60). The buyer's incentive constraint is also unchanged:

$$\xi x'_{a,t} \leq \beta \delta p_a \left(V_{a+1,t+1}^{b'} - V_{0,t+1}^{b'} \right) \iff \xi \lambda^{\sigma-1} x_{a,t} \leq \beta \delta p_a \left(\lambda^{\sigma-1} V_{a+1,t+1}^b - \lambda^{\sigma-1} V_{0,t+1}^b \right). \quad (61)$$

Therefore, any contract that is feasible and incentive compatible with productivity z has a one-to-one transformed counterparty that is feasible and incentive compatible with productivity λz . Moreover, the transformation preserves rankings across contracts. Write the *normalized* optimal contract as the optimal contract in the special case of $z = 1$ as $\{(\tilde{q}_{a,t}, \tilde{x}_{a,t})\}$. Then, for any $z > 0$, the optimal contract policies and corresponding continuation value are given by

$$q_{a,t}(z) = z^\sigma \tilde{q}_{a,t}, \quad x_{a,t}(z) = z^{\sigma-1} \tilde{x}_{a,t}, \quad V_{a,t}^b(z) = z^{\sigma-1} \tilde{V}_{a,t}. \quad (62)$$

It is therefore sufficient to solve once for the normalized optimal contract and scale the solutions appropriately.

Moreover, the same property facilitates dynamic hat algebra. Observe that κ_t enters multiplicatively next to z . We can therefore apply the same arguments to factor out the level of κ_t , so that all that remains in the normalized problem are relative changes in κ_t . In particular, define $\hat{\kappa}_{t,h} \equiv \kappa_{t+h} / \kappa_t$. Substitute $\kappa_{t+h} = \hat{\kappa}_{t,h} \kappa_t$ everywhere. We therefore have

$$q_{a,t}(z) = z^\sigma \kappa_t^{-\sigma} \tilde{q}_{a,t}(\{\hat{\kappa}_{t,h}\}_{h \geq 0}) \quad (63)$$

$$x_{a,t}(z) = z^{\sigma-1} \kappa_t^{1-\sigma} \tilde{x}_{a,t}(\{\hat{\kappa}_{t,h}\}_{h \geq 0}) \quad (64)$$

$$V_{a,t}(z) = z^{\sigma-1} \kappa_t^{1-\sigma} \tilde{V}_{a,t}(\{\hat{\kappa}_{t,h}\}_{h \geq 0}). \quad (65)$$

The same can be done with the demand shifter, ν_t .

B.1.2 Customer capital in [Fitzgerald et al. \(2024\)](#)

Consider a simplified version of the customer-capital accumulation problem in [Fitzgerald et al. \(2024\)](#) (henceforth FHY). A firm with permanent productivity z sells to a given destination. The relevant state is its current customer base, D_t . Conditional on participating in the market, demand

takes the form

$$q_t = Y_t p_t^{-\sigma} D_t^\alpha, \quad (66)$$

where $\sigma > 1$ is the elasticity of demand, $\alpha \in (0, 1)$ governs the importance of customer capital, and Y_t is a market-size shifter. Write the firm's marginal cost as $c_t = \kappa_t/z$.

In FHY's preferred marketing and advertising specification, in which current sales do not directly affect future customer base, the price decision is static and the firm sets the usual CES markup

$$p_t = \frac{\sigma}{\sigma - 1} \frac{\kappa_t}{z}. \quad (67)$$

Substituting the optimal price into operating profits gives

$$(p_t - c_t)q_t = \frac{(\sigma - 1)^{\sigma-1}}{\sigma^\sigma} Y_t \kappa_t^{1-\sigma} z^{\sigma-1} D_t^\alpha \equiv B_t z^{\sigma-1} D_t^\alpha, \quad (68)$$

where B_t collects the aggregate market-size and cost shifters.

The firm can increase its customer base through marketing expenditures, which we denote m_t . Following [Fitzgerald et al. \(2024\)](#), customer base evolves according to

$$D_{t+1} = (1 - \delta)D_t + m_t, \quad (69)$$

and the cost of marketing is $m_t + \varphi \frac{m_t^2}{D_t}$. The firm's post-entry problem is therefore

$$V_t(z, D_t) = \max_{m_t \geq 0} \left\{ B_t z^{\sigma-1} D_t^\alpha - m_t - \varphi \frac{m_t^2}{D_t} + \beta V_{t+1}(z, (1 - \delta)D_t + m_t) \right\}. \quad (70)$$

We exclude sunk costs and fixed costs. Sunk costs can be easily incorporated as the above is the post-entry problem and is therefore unaffected by sunk costs, and a standard [Melitz \(2003\)](#)-Pareto aggregate structure makes entry tractable. Per-period fixed costs may be able to be included under additional assumptions or structure.

Suppose productivity increases from z to λz . Define $r \equiv (\sigma - 1)/(1 - \alpha)$. Consider the following transformation of the state and policy:

$$D'_t = \lambda^r D_t, \quad m'_t = \lambda^r m_t. \quad (71)$$

Observe that the flow payoff scales with λ^r :

$$\pi'_t = B_t (\lambda z)^{\sigma-1} (\lambda^r D_t)^\alpha - \lambda^r m_t - \varphi \frac{(\lambda^r m_t)^2}{\lambda^r D_t} = \lambda^r \pi_t. \quad (72)$$

The law of motion of customer capital also scales:

$$D'_{t+1} = (1 - \delta) \lambda^r D_t + \lambda^r m_t = \lambda^r D_{t+1}. \quad (73)$$

Consider now a normalized policy function $\tilde{m}_t^*(\tilde{D})$, which is the solution to the customer-capital problem in the special case of $z = 1$. This satisfies

$$\tilde{V}_t(\tilde{D}) = \max_{\tilde{m} \geq 0} \left\{ B_t \tilde{D}^\alpha - \tilde{m} - \varphi \frac{\tilde{m}^2}{\tilde{D}} + \beta \tilde{V}_{t+1} \left((1 - \delta) \tilde{D} + \tilde{m} \right) \right\}. \quad (74)$$

For any initial customer base and marketing path, apply the transformation in (71). As shown above, current profits, customer base next period, and hence all future payoffs are simply scaled by z^r . Since rankings across feasible marketing paths are unchanged by the transformation, the transformed version of the $z = 1$ optimal path must be optimal at z . It therefore follows that

$$m_t^*(z, D) = z^r \tilde{m}_t^* \left(\frac{D}{z^r} \right), \quad V_t(z, D) = z^r \tilde{V}_t \left(\frac{D}{z^r} \right). \quad (75)$$

It is therefore sufficient to solve once for a normalized policy function $\tilde{m}_t^*(\tilde{D})$ and scale the solutions appropriately. Unlike in the relational contracting example, the raw state D itself varies with productivity. The productivity-invariant state is instead normalized customer capital, $\tilde{D} = D/z^r$.

As in the previous example, the same property facilitates dynamic hat algebra. The aggregate profitability shifter B_t enters multiplicatively next to $z^{\sigma-1}$. The same scaling argument therefore allows us to factor out its current level, leaving the normalized problem dependent only on relative changes.

B.1.3 Physical capital accumulation

This example proceeds very similarly to the customer capital accumulation case of FHY. Consider a firm with permanent productivity z . The relevant state is its current physical capital stock, k_t . Suppose it chooses investment following a standard capital accumulation problem:

$$V_t(z, k_t) = \max_{i_t} \left\{ A_t z^\eta k_t^\alpha - (\chi/2) i_t^2 / k_t + \beta V_{t+1}(z, (1 - \delta) k_t + i_t) \right\}, \quad (76)$$

with $\eta > 0$ and $\alpha \in (0, 1)$. Unlike the relational contracting problem described above, where the state was relationship age and was invariant to transforming z , here we should expect the capital stock to change with z . The question is whether it is possible to scale productivity together with capital and investment in such a way that a suitably normalized problem retains the same economic trade-offs.

Suppose productivity increases from z to λz . Define $r \equiv \eta/(1 - \alpha)$. Consider the following transformation of the state and policy:

$$k'_t = \lambda^r k_t, \quad i'_t = \lambda^r i_t. \quad (77)$$

Observe that the flow payoff scales with λ^r :

$$\pi'_t = A_t(\lambda z)^\eta (\lambda^r k_t)^\alpha - (\chi/2)(\lambda^r i_t)^2 / (\lambda^r k_t) = \lambda^r \pi_t. \quad (78)$$

The law of motion of capital also scales:

$$k'_{t+1} = (1 - \delta)\lambda^r k_t + \lambda^r i_t = \lambda^r k_{t+1}. \quad (79)$$

Consider now a normalized policy function $\tilde{i}_t(\tilde{k})$, which is the solution to the investment problem in the special case of $z = 1$. This satisfies

$$\tilde{V}_t(\tilde{k}) = \max_{\tilde{i}} \left\{ A_t \tilde{k}^\alpha - (\chi/2) \tilde{i}^2 / \tilde{k} + \beta \tilde{V}_{t+1}((1 - \delta)\tilde{k} + \tilde{i}) \right\}. \quad (80)$$

For any initial capital stock and investment path, consider applying the transformation: $k = z^r \tilde{k}$ and $i = z^r \tilde{i}$. As shown above, current profit and capital next period both scale by z^r . Thus, the entire path and all payoffs are simply scaled by z^r . Since rankings across feasible investment plans are unchanged by the transformation, the transformed version of the $z = 1$ optimal path must be optimal at z . It therefore follows that

$$i_t^*(z, k) = z^r \tilde{i}^* \left(\frac{k}{z^r} \right), \quad V_t(z, k) = z^r \tilde{V}_t \left(\frac{k}{z^r} \right). \quad (81)$$

It is therefore sufficient to solve once for a normalized policy function $\tilde{i}_t^*(\tilde{k})$ and scale the solutions appropriately. As in the relational contracting problem, we can use the same method to factor out the level of A_t , such that the normalized problem depends only on relative changes.

B.1.4 Dynamic screening (Brugués, 2026)

Brugués (2026) studies a partial equilibrium dynamic screening model in which sellers have market power, buyers have persistent private demand types, and the seller offers a dynamic non-stationary menu of contracts. The mechanism must satisfy a continuum of incentive compatibility constraints (it must be optimal for all buyer types to truthfully report their types) and limited enforcement constraints (after reporting their types, the buyers must be willing to honor the contract rather than take the goods and run). Brugués (2026) characterizes the solution and solves the model numerically using the recursive Lagrangean method of Marcet and Marimon (2019).

Let θ denote the buyer's private type, $\hat{\theta}$ denote the type that they report into the mechanism, and $q_t(\hat{\theta})$ and $T_t(\hat{\theta})$ denote the quantity and transfer assigned after reporting $\hat{\theta}$. In the empirical specification used by Brugués (2026), the buyer has isoelastic utility and receives flow payoff

$$u_t^b(\theta, \hat{\theta}) = \theta k q_t(\hat{\theta})^b - T_t(\hat{\theta}), \quad (82)$$

for $b \in (0, 1)$. The seller receive $u_t^s(\hat{\theta}) = T_t(\hat{\theta}) - c_t q_t(\hat{\theta})$.

Suppose that the seller's cost is given by κ_t/z . Suppose that productivity increases from z to λz . Consider a transformation of the mechanism to

$$q'_t(\hat{\theta}) = \lambda^{1/(1-b)} q_t(\hat{\theta}), \quad T'_t(\hat{\theta}) = \lambda^{b/(1-b)} T_t(\hat{\theta}). \quad (83)$$

As in the relational contracting model used in this paper, this re-scaling of the mechanism scales both buyer and seller payoffs by $\lambda^{b/(1-b)}$. The incentive compatibility constraint governing truthful reporting compares lifetime discounted utility under truthful reporting compared to all possible misreports. As these are multiplied by $\lambda^{b/(1-b)}$, if the constraint holds under the original mechanism then it holds under the transformed mechanism. The limited enforcement constraint compares a one-time deviation to the lifetime gains from honoring the contract—again, both are scaled by $\lambda^{b/(1-b)}$ and therefore continue to hold. Moreover, the same argument applies to any combination of misreporting and deviations. As the type distribution does not depend on productivity, the entire feasible set under the original mechanism remains feasible under the transformed mechanism, with the seller's objective scaled. Thus, as in the relational contracting example, it is sufficient to solve for the optimal mechanism in a normalized model, and then scale the solution by productivity.